

Maximizing Miner Revenue in Transaction Fee Mechanism Design*

Ke Wu[†], Elaine Shi[‡] and Hao Chung[§]

Computer Science Department, Carnegie Mellon University

Abstract

Transaction fee mechanism design is a new decentralized mechanism design problem where users bid for space on the blockchain. Several recent works showed that the transaction fee mechanism design fundamentally departs from classical mechanism design. They then systematically explored the mathematical landscape of this new decentralized mechanism design problem in two settings: in the plain setting where no cryptography is employed, and in a cryptography-assisted setting where the rules of the mechanism are enforced by a multi-party computation protocol. Unfortunately, in both settings, prior works showed that if we want the mechanism to incentivize honest behavior for both users as well as miners (possibly colluding with users), then the miner revenue has to be zero. Although adopting a relaxed, approximate notion of incentive compatibility gets around this zero miner-revenue limitation, the scaling of the miner revenue is nonetheless poor.

In this paper, we show that if we make a mild reasonable-world assumption that there are sufficiently many honest users, we can circumvent the known limitations on miner revenue, and design auctions that generate asymptotically optimal miner revenue. We also systematically explore the mathematical landscape of transaction fee mechanism design under the new reasonable-world assumptions, and demonstrate how such assumptions can alter the feasibility and infeasibility landscape.

* Author order is randomized.

[†] kew2@andrew.cmu.edu

[‡] runting@gmail.com

[§] haochung@andrew.cmu.edu

Contents

1	Introduction	2
1.1	Our Results and Contributions	3
1.2	Philosophical Discussions about Our Assumptions and Modeling	6
1.3	Additional Related Work	7
2	Technical Roadmap	8
2.1	Transaction Fee Mechanism	8
2.2	Infinite Block Setting	9
2.3	Finite Block Setting	13
2.3.1	Strict Incentive Compatibility	13
2.3.2	Approximate Incentive Compatibility	14
2.4	Additional Results	15
3	Model and Definitions	15
3.1	MPC-Assisted Model	16
3.2	Defining Incentive Compatibility	17
3.2.1	Bayesian Incentive Compatibility	18
3.2.2	Ex Post Incentive Compatibility	19
4	Feasibility for Infinite Block Size	19
4.1	MPC-Assisted, Threshold-Based Mechanism	19
4.2	Analysis of the LP-Based Mechanism	20
4.2.1	Preliminaries: Linear Algebra Tools	20
4.2.2	Proofs for the LP-Based Mechanism	21
5	Characterization for Finite Block Size	27
5.1	Characterization for Strict IC	27
5.1.1	Feasibility for $c = 1$	27
5.1.2	Zero Social Welfare for Users When $c \geq 2$	28
5.2	Feasibility for Approximate IC: Diluted Threshold-Based Mechanism	32
6	Bounds on Miner Revenue	33
6.1	Known- h Model	34
6.2	Necessity of Bayesian Incentive Compatibility	36
6.3	Honest Majority of Bids	38
A	MPC-Assisted, Parity-Based Mechanism	43

1 Introduction

The transaction fee mechanism (TFM) [LSZ19, Yao, BEOS19, BCD⁺, Rou20, Rou21, FMPS21, CS23, GY22, ZCZ22] is a new decentralized mechanism design problem that arises in a blockchain environment. Since the space on the blockchain is scarce, users must bid to get their transactions included and confirmed whenever a new block is minted. Earlier works observed that mechanism design in a decentralized environment departs fundamentally from classical mechanism design [LSZ19, Yao, BEOS19, BCD⁺, Rou20, Rou21, FMPS21, CS23, GY22, ZCZ22]. The majority of classical mechanisms assume that the auctioneer is trusted and will honestly implement the prescribed mechanism. However, in a decentralized environment, the auction is implemented by a set of miners or consensus nodes¹ who are incentivized to take advantage of profitable deviations (if there are any) rather than implementing the mechanism honestly. As a simple example, while the Vickrey auction [Vic61] (a.k.a., the second-price auction) is considered an awesome auction by classical standards, it is not a great fit for a decentralized environment as explained below. Suppose we confirm k bids and they all pay the $(k + 1)$ -th price to the miner. Then the miner could inject a fake bid that is slightly less than the k -th price, thus causing confirmed bids to pay essentially the k -th price. Alternatively, the same effect can also be achieved if the miner colludes with the $(k + 1)$ -th bidder and asks it to raise its bid to almost exactly the k -th price. The coalition can then split off the additional gains off the table, using *binding* side contracts that can be instantiated through the decentralized smart contracts that are available in blockchain environments.

This drives us to rethink what is a “dream” TFM in the absence of a fully trusted auctioneer. Recent works [BCD⁺, Rou20, Rou21, CS23] formulated the following desiderata:

- *User incentive compatibility (UIC)*: a user’s best strategy is to bid truthfully, even when the user has observed others’ bids.
- *Miner incentive compatibility (MIC)*: the miner’s best strategy is to implement the mechanism honestly, even when the miner has observed all users’ bids.
- *Side-contract-proofness (SCP)*: playing honestly maximizes the joint utility of a coalition consisting of the miner and at most c users, even after having observed all others’ bids.²

Roughgarden showed that Ethereum’s EIP-1559 mechanism can simultaneously achieve all three properties, as long as the block size is *infinite* [BCD⁺, Rou20]. In practice, EIP-1559 tries to be in the “infinite block size” regime by estimating a reserve price based on the recent history. The reserve price is a minimum threshold used to filter the transactions to avoid congestion. For the case of finite block size (i.e., when congestions do occur), Chung and Shi [CS23] showed that it is impossible to satisfy all three properties at the same time without the use of cryptography. However, the subsequent work of Shi, Chung, Wu [SCW23] shows that we can use the *MPC-assisted model* to circumvent this impossibility. In the MPC-assisted model, the TFM’s rules are securely enforced by a multi-party computation protocol among the miners, thus taking away the miner’s ability to unilaterally decide which transactions to include in the block. Variants of the MPC-assisted model are being developed by mainstream blockchain projects such as Ethereum. In particular, the community has been making an effort to build “encrypted mempools” [enc], which can be viewed as a concrete instantiation of the MPC-assisted model³. Under the MPC-assisted

¹Throughout the paper, we call the consensus nodes “miners” regardless of whether the consensus protocol uses proof-of-work or proof-of-stake.

²The formal definition of joint utility is given in Section 3.1.

³Exactly whether encrypted mempool realizes the MPC ideal functionality needed by Shi, Chung, Wu [SCW23] requires a rigorous proof.

model, [SCW23] showed that one can indeed construct a TFM that simultaneously satisfies UIC, MIC, and SCP (for $c = 1$) under finite block size.

Limit on miner revenue. Unfortunately, no matter whether in the plain or in the MPC-assisted model, all these prior works [Rou20, CS23, SCW23] suffer from a “zero-miner-revenue limitation”: all the payment from the users is burnt⁴ and the miner obtains zero revenue. The works of Chung and Shi [CS23] and Shi, Chung, Wu [SCW23] proved that this limitation is in fact inherent. Specifically, any TFM (either in the plain model or MPC-assisted model) that simultaneously satisfies UIC, MIC, and SCP (even for $c = 1$) must have zero miner revenue. In both the plain and the MPC-assisted models, the zero miner-revenue limitation holds in a very strong sense: regardless of whether the block size is finite or infinite, and even when the miner colludes with at most $c = 1$ user.

Moreover, [SCW23] additionally explored whether relaxing the *strict* incentive compatibility notion to *approximate* incentive compatibility can increase the miner revenue. They show a somewhat pessimistic, *unscalability of miner revenue* result. Specifically, with ϵ -incentive compatibility, the miner revenue cannot enjoy linear scaling w.r.t. the magnitude of the bids. Given the landscape, we ask the following natural question:

Can we circumvent the severe limitation on miner revenue, under some reasonable-world assumptions?

1.1 Our Results and Contributions

We are inspired by the philosophy adopted by a line of work at the intersection of cryptography and game theory [HT04, KN08, ADGH06, OPRV09, AL11, ACH11, GKM⁺13, GKTZ15, GTZ15, Kat08, DR07, GLR10, CGL⁺18, WAS22, CCWS21, PS17, KMSW22, FW20, EFW22]. In these works, the game theoretic properties hold as long as sufficiently many players are honest. Because the game theoretic guarantees ensure that honest behavior is an equilibrium, and that players are incentivized to behave honestly, this in turn reinforces the “sufficient honesty” assumption.

Therefore, we ask whether we can overcome the severe limitation on miner revenue also under some type of “sufficient honesty” assumption. Phrasing the precise “sufficient honesty” assumption, however, turns out to be technically subtle, partly because TFMs must work in an *open* setting where anyone can post a bid, and the mechanism is unaware of the number of bids a-priori. One naïve attempt is to assume that among the bids posted, half of them come from honest users. Unfortunately, this approach does not work. In Section 6.3, we show that even under such an “honest majority bids” assumption, we would still suffer from an $O(1)$ -miner revenue limitation.

Reasonable-world assumption: known lower bound on the number of honest users. Instead of the “honest majority bids” assumption, we make a subtly different assumption — we assume that there is an a-priori known lower bound h on the number of honest users. Note that this assumption also promises that at least h users will show up. We refer to this as the *known- h model*. In this model, we first observe that the zero miner-revenue limitation no longer holds. Instead, we can prove an $O(h)$ -limit on the miner revenue as stated in the following theorem.

Theorem 1.1 (Informal: limit on miner revenue in the known- h model). *In the known- h model, no MPC-assisted mechanism that simultaneously satisfies UIC, MIC, and SCP (even in the Bayesian*

⁴We stress that the miners can still get a fixed block reward which incentivizes them to mine — in fact, Ethereum’s EIP-1559 burns all base fees and pays the miner a fixed block reward. This fixed block reward does not affect the game-theoretic analysis and thus is typically ignored in the game-theoretic modeling [Rou20, CS23, SCW23].

setting) can achieve more than $h \cdot \mathbf{E}(\mathcal{D})$ expected miner revenue where $\mathbf{E}(\mathcal{D})$ denotes the expectation of the value distribution \mathcal{D} .

More generally, in the known- h model, if the number of users is n , no MPC-assisted mechanism that simultaneously satisfies ϵ -UIC, ϵ -MIC, and ϵ -SCP (even in the Bayesian setting) can achieve more than $h \cdot \mathbf{E}(\mathcal{D}) + \frac{2(n-h)}{\rho} (\epsilon + C_{\mathcal{D}}\sqrt{\epsilon})$ expected miner revenue, where ρ is an upper bound on the fraction of miners controlled by the strategic coalition, and $C_{\mathcal{D}} = \mathbf{E}_{X \sim \mathcal{D}}[\sqrt{X}]$.

Furthermore, the above limitation holds no matter when the block size is finite or infinite, and even when the miners collude with at most $c = 1$ user.

In the above theorem, $\epsilon \geq 0$ is a parameter that measures the slack in the incentive compatibility notion. When $\epsilon = 0$, there is no slack, and we achieve strict incentive compatibility. One informal interpretation of the above theorem is the following: for ϵ incentive compatibility, Theorem 1.1 allows us to hope for a mechanism where roughly speaking, from each of h users, the miners can hope to get $\mathbf{E}(\mathcal{D})$ revenue which scales proportionally w.r.t. to the bid distribution \mathcal{D} . For each of the remaining users, the miners can potentially get some function that depends on ϵ and the bid distribution \mathcal{D} , but the term does not scale linearly w.r.t. the magnitude of the bid distribution for natural distributions.

The above Theorem 1.1 allows us to hope for a TFM in the known- h model that achieves revenue that scales with h as well as the magnitude of the bid distribution \mathcal{D} . So can we indeed design a mechanism with asymptotically optimal mine revenue matching Theorem 1.1?

Mechanisms for infinite block size. For the infinite block size regime, we propose two mechanisms in the MPC-assisted model:

- The first one, called *threshold-based mechanism*, is a simple and practical mechanism that satisfies almost-strict incentive compatibility except for a tiny slack ϵ that is exponentially small in h .
- The second one, called *LP-based mechanism* (since it uses linear programming), is a result of theoretical interest. It achieves *strict* incentive compatibility, but under one extra assumption (besides a-priori knowledge of h), that the number of fake bids injected by the strategic coalition is bounded.

Both mechanisms achieve asymptotically optimal miner revenue⁵ w.r.t. Theorem 1.1. We assume that only honest users' true values are i.i.d. sampled from some distribution \mathcal{D} , whereas the strategic users' true values can be *arbitrary* non-negative real numbers. Next, we state the corresponding theorems for the two mechanisms below:

Theorem 1.2 (Informal: threshold-based mechanism). *Suppose that honest users' values are sampled i.i.d. from some distribution \mathcal{D} . Then, there exists an MPC-assisted TFM in the known- h model that satisfies ex post UIC, Bayesian ϵ -MIC, and Bayesian ϵ -SCP (for any number of colluding users) for $\epsilon = O_{\mathcal{D}}(\exp(-\Omega(h)))$ where $O_{\mathcal{D}}(\cdot)$ hides terms related to the value distribution \mathcal{D} . Furthermore, the expected total miner revenue $\Theta(h) \cdot \text{median}(\mathcal{D})$.*

Essentially, from each of h users, the miners can obtain revenue that scales linearly w.r.t. both the bid magnitude. By contrast, without the known- h assumption, for our choice of ϵ which is exponentially small in h , the miner revenue must be exponentially small in h as shown in prior work [SCW23]. Observe also that the miner revenue is asymptotically optimal up to additive factors that are exponentially small in h due to Theorem 1.1.

⁵We achieve asymptotic optimal miner revenue w.r.t. h assuming that the expectation and median of the distribution \mathcal{D} is a constant independent of h .

Theorem 1.3 (Informal: LP-based mechanism). *Suppose that honest users’ values are sampled i.i.d. from some distribution \mathcal{D} . Then, there exists an MPC-assisted TFM that satisfies ex post UIC, Bayesian MIC, and Bayesian SCP where the MIC and SCP guarantees hold as long as the total number of bids d contributed by the strategic coalition satisfies $d \leq \frac{1}{8} \sqrt{\frac{h}{2 \log h}}$. Further, the expected total miner revenue is $\Theta(h) \cdot \text{median}(\mathcal{D})$.*

In the above theorem, we need the extra assumption that the strategic coalition does not control too many bids. Effectively, this is assuming that the coalition cannot inject too many fake bids. Currently, we do not know whether this extra assumption (besides known- h) is needed to overcome the zero miner-revenue limitation. We leave this as an interesting open question.

Mechanisms for finite block size. For the finite block size case, we propose two mechanisms:

- We propose a simple mechanism called *diluted threshold-based mechanism* that achieves *approximate* incentive compatibility. Further, for sufficiently large h , the mechanism achieves asymptotically optimal miner revenue.
- For theoretical interest, we propose another mechanism called *LP-based mechanism with random selection* which achieves *strict* incentive compatibility and asymptotically optimal miner revenue — but under the additional assumptions that the coalition cannot inject too many fake bids, and moreover, the miners collude with at most $c = 1$ user. Jumping ahead, the $c = 1$ assumption will be later justified in Theorem 1.6.

We state the corresponding theorems for the two mechanisms below:

Theorem 1.4 (Informal: diluted threshold-based mechanism). *Suppose the block size is k , and that honest users’ values are sampled i.i.d. from some bounded distribution \mathcal{D} . Then, there exists an MPC-assisted TFM in the known- h model that satisfies ex post UIC, Bayesian ϵ -MIC, and Bayesian ϵ -SCP (for any number of colluding users) for $\epsilon = O_{\mathcal{D}}(\exp(-\Omega(h)))$. Furthermore, for sufficiently large h , the mechanism achieves expected total miner revenue $\Theta(k) \cdot \text{median}(\mathcal{D})$.*

Theorem 1.5 (Informal: LP-based mechanism with random selection). *Suppose the block size is k , and suppose that honest users’ values are sampled i.i.d. from some distribution \mathcal{D} . Then, there exists an MPC-assisted TFM that satisfies ex post UIC, Bayesian MIC, and Bayesian SCP, where the MIC and SCP guarantees hold when 1) at most $c = 1$ user colludes when miners, and 2) the total number of bids d contributed by the strategic coalition satisfies $d \leq \frac{1}{8} \sqrt{\frac{h}{2 \log h}}$. Further, the expected total miner revenue is $\Theta(\min\{h, k\}) \cdot \text{median}(\mathcal{D})$.*

We justify the $c = 1$ assumption in the LP-based mechanism with random selection by proving the following impossibility result: for finite block size, no “interesting” mechanism can simultaneously achieve UIC, MIC, and SCP for $c \geq 2$ even in the MPC-assisted model. Specifically,

Theorem 1.6 (Informal: finite block, $c \geq 2$). *Even in the known- h model, any MPC-assisted TFM that simultaneously satisfies Bayesian UIC, Bayesian MIC, and Bayesian SCP for $c \geq 2$ must suffer from 0 expected social welfare for the users under a bid vector $\mathbf{b} \sim \mathcal{D}^\ell$ where $\ell > h$.*

Necessity of Bayesian equilibrium. All of our feasibility results, namely, Theorems 1.2 to 1.5, rely on a Bayesian notion of equilibrium (for the MIC and SCP guarantees). As argued by [SCW23], the Bayesian notion of equilibrium is suitable for the MPC-assisted model since the users cannot observe others’ bids before submitting their own.

We show that the reliance on Bayesian notions of equilibrium is necessary (see Section 6.2) — had we insisted on an *ex post* notion of equilibrium in the MPC-assisted model, our additional reasonable-world assumptions would not help us overcome the previously known impossibility results. More specifically, we show that any MPC-assisted mechanism that simultaneously achieves ex post UIC and SCP must suffer from zero miner revenue even in the known- h model. Similarly, for approximate but ex post notions of incentive compatibility, the same miner revenue limitation stated in [SCW23] still applies even in the known- h model. Further, the above restrictions on miner revenue hold no matter whether the block size is finite or infinite.

1.2 Philosophical Discussions about Our Assumptions and Modeling

Known- h assumption. Our assumption about a known upper bound h on the number of honest users has the following justifications:

- First, as mentioned, the zero miner-revenue limitation in earlier works holds in a very strong sense, and even when we make a cryptographic style assumptions such as “a majority of the users or bids are honest” — see Section 6.3 for more details.
- Second, TFMs must work in an *open setting* where anyone can post a bid and the mechanism or players do not know the number of bids in advance. This is also an important reason why TFMs depart from classical mechanism design. The precise assumption we need for circumventing the zero miner revenue limitation is an absolute lower bound h on the number of honest users. Simply assuming that the majority of the bids are honest is not sufficient (see Section 6.3).
- Finally, our definitional framework ensures that *honest behavior is an equilibrium*, and thus players are incentivized to behave honestly. This, in turn, reinforces the h -honest users assumption (as long as enough users show up). As mentioned earlier, the same philosophy has been adopted in a line of prior works at the intersection of game theory and cryptography [HT04, KN08, ADGH06, OPRV09, AL11, ACH11, GKM⁺13, GKTZ15, GTZ15, Kat08, DR07, GLR10, CGL⁺18, WAS22, CCWS21, PS17, KMSW22, FW20, EFW22].

Limited fake bids. Recall that besides the known- h assumption, our results that achieve strict incentive compatibility require an extra assumption that the number of fake bids injected by the coalition is limited. This assumption is motivated and justified by the following observations. To submit fake bids, the strategic player or coalition needs to have some coin or account with a non-zero balance. Given that the strategic player has a limited initial budget, it cannot control infinitely many accounts. Moreover, if one posts multiple conflicting transactions double-spending the same coin or units of currency, they can easily be detected and suppressed.

As mentioned, we currently do not know whether this limited fake bids assumption can be removed while still achieving strict incentive compatibility. We pose this as an open question.

Robustness w.r.t. parameter estimation. Among our proposed mechanisms, the ones that achieve approximate incentive compatibility, namely, threshold-based or diluted threshold-based mechanisms are simpler and more practical. Just like how Ethereum’s EIP-1559 needs to estimate a suitable base fee, these mechanisms also need to estimate some parameters a-priori. In particular, our (diluted) threshold-based mechanism needs to know an estimate of h and the median of the value distribution \mathcal{D} in advance. Just like Ethereum’s EIP-1559, we can estimate these parameters from recent history. For example, one can estimate the total number of bids n from the degree of congestion observed in recent blocks. Now, if we are willing to assume that half of the n anticipated

bids are honest (note that our mechanisms incentivize honest behavior), we can get an estimate of h . Similarly, one can estimate the median of the distribution \mathcal{D} from the recent history too.

One important observation is that our threshold-based mechanism and diluted threshold-based mechanisms are quite robust to errors in the estimates. As mentioned later in Remark 2.2, if we set the threshold to $\hat{h}/4$ for some estimated \hat{h} , and let \hat{m} be the estimated median, then the mechanisms will achieve approximate incentive compatibility as long as $h_{\text{real}} \cdot q_{\text{real}} \geq \hat{h} \cdot \frac{(1+\delta)}{4}$ for some arbitrarily small constant $\delta > 0$, where h_{real} is the actual number of honest users, and q_{real} is the actual percentile of the estimate \hat{m} . For example, if $h_{\text{real}} = 0.6h$, and $q_{\text{real}} = 40\%$, then our mechanisms still satisfy approximate incentive compatibility for an exponentially small ϵ .

Independent identically distributed true values assumption. All the mechanisms in our paper only assume that honest users’ true values are i.i.d. sampled from the distribution \mathcal{D} . The strategic users can have arbitrary non-negative true values.

1.3 Additional Related Work

We now review some closely related recent works besides the prior works on transaction mechanism design [LSZ19, Yao, BEOS19, BCD⁺, Rou20, Rou21, FMPS21, CS23] already mentioned.

TFM in a Bayesian setting. The recent works of Zhao, Chen, and Zhou [ZCZ22] and Gafni and Yaish [GY22] both consider TFM in a Bayesian setting. Although their works did not explicitly define the MPC-assisted model, from a practical standpoint, their results are in fact only relevant in an MPC-assisted (or a similar) model. As explained in Section 3.2 and Fact 3.3, plain-model TFMs that achieve *Bayesian* equilibrium also achieve *ex post* equilibrium, since in the plain-model game, the strategic player can decide its actions *after* having observed honest users’ bids.

Gafni and Yaish [GY22] suggest a mechanism that satisfies Bayesian UIC, while also satisfying MIC and OCA-proof (short for offchain-agreement-proof) even if the miner knows everyone’s bid. Further, their mechanism works in the finite-block setting while achieving asymptotical optimality in social welfare and revenue. We stress that their result does not contradict the zero miner-revenue limitation proven by [SCW23] since their OCA-proofness notion (originally defined by Roughgarden [Rou20, Rou21]) is of a different nature from our side-contract-proofness (SCP) notion (originally defined by Chung and Shi [CS23]). Roughly speaking, OCA-proofness requires that a strategic coalition cannot enter an off-chain contract that increases *everyone’s* utility (*not just those in the coalition*) relative to what’s achievable on-chain. In comparison, SCP is the notion that directly captures the cryptocurrency community’s outpouring concerns about Miner Extractable Value (MEV). In particular, middleman platforms such as Flashbot facilitate the collusion of miners and users, where the coalition plays strategically to profit themselves at the expense of other users. This is why we choose to use the SCP notion rather than OCA-proofness. Moreover, the reason why the cryptocurrency community is developing encrypted mempool techniques (which can be viewed as instantiations of the MPC-assisted model) is also because they care about SCP (i.e., resilience to MEV).

Zhao, Chen, and Zhou [ZCZ22] suggest a mechanism that generates positive miner revenue while achieving Bayesian UIC and Bayesian 1-SCP even for the finite block setting. Their result does not contradict the 0-miner revenue limitation of [SCW23], since Zhao, Chen, Zhou [ZCZ22] consider only a restricted strategy space. In their work, a strategic user or a miner-user coalition can only deviate by bidding untruthfully; the coalition cannot inject fake bids, strategic users cannot drop out, and nor can strategic miners alter the inclusion rule. Due to their restricted strategy space, their results are only relevant under very stringent assumptions: 1) the TFM is implemented in the

MPC-assisted (or similar) model; 2) the TFM is fully “permissioned” and allows only a set of pre-registered users to submit bids. In particular, the latter “permissioned” requirement is unrealistic for major decentralized cryptocurrencies today where any user can join and submit transactions.

Cryptography meets game theory. Prior to the advent of cryptocurrencies, a line of work [HT04, KN08, ADGH06, OPRV09, AL11, ACH11, GKM⁺13, GKTZ15, GTZ15, Kat08, DR07, GLR10, CGL⁺18, WAS22, CCWS21, PS17, KMSW22, FW20, EFW22] investigated how cryptography and game theory can help each other. For example, cryptography can help remove the trusted mediator assumption in correlated equilibria [DR07]. Adopting game-theoretic fairness can allow us to circumvent lower bounds pertaining to the more stringent cryptographic notions of fairness [HT04, ADGH06, IML05, OPRV09, CGL⁺18, WAS22]. Ferreira and Weinberg [FW20] and Essaidi, Ferreira and Weinberg [EFW22] showed that cryptographic commitments can help us circumvent impossibilities pertaining to credible auctions. As Chung and Shi [CS23] explained in detail, credible auction is of a fundamentally different nature from transaction fee mechanism design.

2 Technical Roadmap

2.1 Transaction Fee Mechanism

Environment. For preciseness in our subsequent formal description, we define an (h, ρ, c, d) -environment, where h is the promised lower bound on the number of honest users, $\rho \in [0, 1]$ is the fraction of strategic miners, c is the maximum number of strategic users that collude with miners, and d is the maximum number of bids contributed by the strategic coalition, i.e. fake bids.

Universality. We can replace a subset of these variables with a wildcard $*$ if the mechanism achieves incentive compatibility no matter what the variable turns out to be. For example, a TFM that achieves incentive compatibility in an $(*, \rho, c, *)$ -environment if it works when the maximum fraction of strategic miners is ρ , and the maximum number of colluding users is at most c , and regardless of how many honest users there are and how many bids are contributed by strategic individuals or the coalition. In this case, we also say that the mechanism is universal in the parameters h and d . Using this notation, the mechanisms described by Shi, Chung, Wu [SCW23] are universal in the parameters h and d . Similarly, the limitation on miner revenue they prove can also be interpreted as a limitation of mechanisms that are universal in h and d .

Transaction fee mechanism in the MPC-assisted model. As in previous works, we consider a single auction instance that decides which transactions can be confirmed in the next block. In this paper, “bids” and “transactions” are used interchangeably. A transaction fee mechanism (TFM) in the MPC-assisted model consists of the following randomized algorithms.

- *Confirmation rule* chooses a subset of at most k bids to confirm, where k denotes the block size.⁶
- *Payment rule* decides how much each confirmed bid pays.
- *Miner revenue rule* decides how much revenue the miners get.

⁶In the MPC-assisted model, the mechanism is implemented by the ideal functionality, and the miners cannot decide which transactions are included and considered as the input of the mechanism. Since all mechanisms proposed in our paper work in the MPC-assisted model, we simplify the definition of TFM compared to [CS23], where the TFM in [CS23] also needs to specify the *inclusion rule* that tells the miner which transactions are included in the next block. However, a strategic miner can still inject some fake bids. The formal strategy space is defined in Section 3.

We consider a single-parameter environment, i.e., each user has a transaction with the true value represented by a single, non-negative real number. A user’s utility equals its true value minus its payment if its transaction gets confirmed in the block. An honest user submits one bid $b \in \mathbb{R}$ representing its true value, while an honest miner submits zero bids and follows the protocol honestly. Strategic users or miners, however, can choose to register one or more identities and submit arbitrary bids under these identities. Moreover, they can choose to drop out during the execution. When a strategic user or miner submits multiple bids, they can only obtain the value when the bid with the true value is confirmed, and other bids are considered fake bids. Fake bids have no intrinsic value to the users and the miners even when they are confirmed.

All mechanisms proposed in our paper work in the MPC-assisted model. We consider the same, generic MPC-assisted model as [SCW23], where the miners jointly execute an MPC protocol that securely evaluates the rules of the TFM. Further, the miners equally divide up the revenue among themselves. A formal description of the model can be found in Section 3.

To illustrate the technical highlights, it is convenient to abstract out the multi-party computation protocol as an ideal functionality. Therefore, we can think of the game as follows:

- Each player (either user or miner) first submit zero to multiple bids to the ideal functionality.
- The ideal functionality then executes the rules of the TFM, outputs the set of confirmed bids, how much each bid needs to pay, and how much the miners get based on the prescribed rules.

Notably, in the MPC-assisted model, players cannot see others’ bids when posting their own. In particular, in an actual instantiation of the above functionality, the players would send bids either in a secret-shared or encrypted format. For this reason, Bayesian notions of equilibrium make sense in the MPC-assisted model. [SCW23] suggested one way to securely realize the above ideal functionality. Meanwhile, efforts to build an “encrypted mempool” by the cryptocurrency community can also be viewed an alternative way to instantiate the ideal functionality — although whether the suggested approaches provably realize this ideal functionality is yet to be proven.

[SCW23] argued that as a starting point, it makes sense to first explore such a generic functionality, since currently, we lack understanding how cryptography in general can help decentralized mechanism design from a game theoretic perspective. Once we understand this better, we can then focus on making customized optimizations for the actual computational tasks that need to be securely evaluated. Our work is motivated by the same philosophy, i.e., we focus on understanding the game theoretic aspects rather than concrete optimizations for realizing the MPC.

Universality in ρ in the idealized model. As mentioned, it is convenient to think of the MPC as an ideal functionality. In this idealized model, our mechanisms can achieve universality in ρ . However, when the ideal functionality is actually instantiated, e.g., with an honest-majority MPC protocol, the resulting protocol would require $\rho < 50\%$.

2.2 Infinite Block Setting

For the infinite block setting, we can achieve $\Theta(h)$ miner revenue in (h, ρ, c, d) -environments. To aid understanding, we first present a simple parity-based mechanism that works for $h = 1$, and then we present our main results.

Glimpse of hope. First, consider the special case where we are promised that there is at least $h = 1$ honest user. In this case, the following simple parity-based mechanism satisfies ex-post UIC, Bayesian MIC, and Bayesian SCP in $(1, *, *, *)$ -environments.

MPC-assisted, parity-based mechanism

// Let m be the median of the distribution \mathcal{D} such that $\Pr_{x \sim \mathcal{D}}[x \geq m] = 1/2$.

- All bids that are at least m get confirmed and pay m .
- If the number of confirmed bids is odd, then the total miner revenue is m ; else the total miner revenue is 0.

In the above mechanism, as long as there is at least one honest bid, the expected miner revenue is always $m/2$ no matter how the coalition behaves. This is because the strategic coalition cannot predict whether the honest bid is bidding at least the median or not. With this key observation, it is not hard to see that the mechanism satisfies Bayesian MIC and Bayesian SCP (for an arbitrary c). Further, ex post UIC follows directly since the mechanism is a simple posted-price auction from a user's perspective.

Observe also the following subtlety: when the number of confirmed bids is odd, it implies that there is at least one confirmed bid. Therefore, the mechanism guarantees that the miner revenue does not exceed the total payment.

Remark 2.1 (A note about the median assumption). In the above, we assumed that the bid distribution \mathcal{D} has a median m such that $\Pr_{x \sim \mathcal{D}}[x \geq m] = 1/2$. In case the median m does not exactly equally divide the probability mass half and half, then it must be that $\Pr_{x \sim \mathcal{D}}[x > m] < 1/2$ and $\Pr_{x \sim \mathcal{D}}[x < m] < 1/2$. In this case, we can modify the above mechanism slightly as follows: if a user's bid is strictly greater than m , then it is confirmed; if a user's bid is exactly m , then we confirm it with some appropriate probability q ; else the user's bid is not confirmed. We can always pick a q such that a bid randomly sampled from \mathcal{D} is confirmed with probability exactly $1/2$. Finally, the miner revenue rule is still decided the same way as before.

Threshold-based mechanism. The parity-based mechanism overcomes the 0 miner-revenue limitation by assuming the existence of at least $h = 1$ honest user. However, the drawback is obvious: the total miner revenue is severely restricted and does not increase w.r.t. the number of bids. A natural question is whether we can achieve $O(h)$ expected miner revenue for general h .

We give an affirmative answer. We first present a simple, practical mechanism called the threshold-based mechanism that achieves almost-strict incentive compatibility except for a tiny slack ϵ that is exponentially small in h . Then, for theoretical interest, we present another mechanism that achieves strict incentive compatibility but requires an extra assumption on the number of bids contributed by strategic players.

MPC-assisted, threshold-based mechanism

// Let m be the median of the distribution \mathcal{D} such that $\Pr_{x \sim \mathcal{D}}[x \geq m] = 1/2$.

- All bids that are at least m get confirmed and pay m .
- If the number of confirmed bids is at least $h/4$, then the miner revenue is $m \cdot h/4$; else the total miner revenue is 0.

Due to the standard Chernoff bound, except with $e^{-\Omega(h)}$ probability, the number of confirmed bids among the h (or more) honest bids is at least $h/4$. Therefore, the above mechanism achieves at least $m \cdot h/4 \cdot (1 - e^{-\Omega(h)})$ expected miner revenue. If the number of confirmed honest bids is $h/4$ or higher, then the coalition cannot increase the miner revenue no matter how it behaves. Only when the number of confirmed honest bids is less than $h/4$, is it possible for the coalition to influence the miner revenue by at most $m \cdot h/4$. Therefore, it is not hard to see that the mechanism satisfies ϵ -Bayesian MIC and ϵ -Bayesian SCP in $(h, *, *, *)$ -environments, for $\epsilon = m \cdot h \cdot e^{-\Omega(h)}/4$.

Just like before, in case the median m does not exactly divide the probability mass half and half, we can use the same approach of Remark 2.1 to modify the mechanism and make it work.

Remark 2.2 (On the robustness of parameter estimation). The threshold-based mechanism requires the mechanism to estimate h and the median m of the distribution \mathcal{D} . Just like how Ethereum EIP-1559 estimates its base price, we can estimate h and the median from recent history. In particular, from the congestion level in recent blocks, we can estimate a lower bound on the total number of bids. Now, assuming that at least half of them are honest (recall that our mechanism incentivizes honesty), we can get an estimate of h correspondingly. Similarly, we can estimate the median of the bid distribution from past history.

An advantage of the threshold-based mechanism is that it is quite tolerant of errors in the estimation. For example, if the estimated m is actually the 40-percentile of \mathcal{D} , and the actual number of honest users is only $0.7h$ where h is the estimate used by the mechanism, the expected number of users bidding at least m is at least $0.4 \cdot 0.7h = 0.28h$. In this case, we can still guarantee that except with exponentially small in h probability, at least $h/4$ users will bid at least m . Thus, the resulting mechanism would still be almost strictly incentive compatible except for a slack ϵ that is exponentially small in h .

LP-based mechanism. Although the ϵ slack in the threshold-based mechanism is exponentially small which is not a problem in practice, it is still theoretically interesting to ask whether we can get $\Theta(h)$ total miner revenue but with *strict* incentive compatibility. To achieve this, our idea is to devise a mechanism that is “close in distance” to the aforementioned threshold-based mechanism, but correcting the “error” such that we can achieve strict incentive compatibility.

Observe that the earlier threshold-based mechanism only needs an a-priori known lower bound on h , and it is universal in the parameters c and d . To achieve strict incentive compatibility, we additionally assume that the the number of bids contributed by the strategic coalition is upper bounded by some a-priori known parameter d .

Now, consider the following mechanism that relies on linear programming to correct the error in the earlier threshold-based mechanism. For simplicity, we assume that the honest bid distribution \mathcal{D} has a median m such that $\Pr_{x \sim \mathcal{D}}[x \geq m] = 1/2$ — if not, we can again use the technique of Remark 2.1 to modify the mechanism and make it work.

MPC-assisted, LP-based mechanism

// Let m be the median of the distribution \mathcal{D} , i.e., $\Pr_{x \sim \mathcal{D}}[x \geq m] = 1/2$.

- All bids that are at least m get confirmed and pay m .
- Let n be the length of the bid vector, let $\mathbf{y} := (y_0, y_1, \dots, y_n)$ be any feasible solution to the following linear program:

$$\forall i \in [n] : 0 \leq y_i \leq i \cdot m \tag{1}$$

$$\forall 0 \leq j \leq d : \sum_{i=0}^{n-d} q_i \cdot y_{i+j} = \frac{m \cdot h}{4} \tag{2}$$

where q_i is the probability of observing i heads if we flip $n - d$ independent fair coins.

- The total miner revenue is y_s where s is the number of bids confirmed.

In the above, Equation (1) expresses a *budget feasibility* requirement, i.e., the total miner revenue cannot exceed the total user payment. Equation (2) expresses a *fixed-revenue requirement*

stipulating that the miner revenue must be exactly $m \cdot h/4$ no matter how the strategic individual or coalition behaves (as long as it controls at most d bids). More specifically, Equation (2) contains one requirement for each $j \in [0, d]$: conditioned on the fact that among the (at most) d bids controlled by the strategic individual or coalition, exactly j of them are confirmed, the expected miner revenue must be exactly $m \cdot h/4$ where h is an a-priori known lower bound on the number of honest users.

Remark 2.3. We know that the actual number of honest users that show up is at least $\max(n - d, h)$. So if $n - d > h$, it means that more honest users showed up than the anticipated number h . Observe that on the left-hand side of Equation (2), we are tossing coins for $n - d$ honest users' bids. However, it is important that the right-hand-side of Equation (2) use the a-priori known h rather than the observed $n - d$; otherwise, injecting extra (but up to $d - c$) fake 0-bids can increase the expected miner revenue, which violates MIC and SCP.

If the LP in the above mechanism indeed has a feasible solution, then we can prove that the resulting mechanism satisfies ex post UIC, Bayesian MIC, and Bayesian SCP in $(h, *, *, d)$ -environments. The formal proofs are presented in Section 4.2.2.

The key technical challenge is to answer the question of why the LP has a feasible solution. Intuitively, the earlier threshold-based mechanism gives an ‘‘approximate’’ solution $\hat{\mathbf{y}} := (\hat{y}_0, \dots, \hat{y}_n)$ to the LP, where $\hat{y}_i = 0$ for $i \leq (n - d)/4$ and $\hat{y}_i = \frac{m \cdot h}{4}$ otherwise. With the approximate solution $\hat{\mathbf{y}}$, the equality constraints in Equation (2) may be satisfied with some small error. We want to show that we can adjust the $\hat{\mathbf{y}} := (y_0, y_1, \dots, y_n)$ vector slightly such that we can correct the error, and yet without violating the budget feasibility constraints (Equation (1)).

To achieve this, we will take a constructive approach. We first guess that a feasible solution is of the form $\mathbf{y} = \hat{\mathbf{y}} + \mathbf{e}$ where \mathbf{e} is a correction vector that is zero everywhere except in the coordinates $\tau, \tau + 1, \dots, \tau + d$ for some appropriate choice of τ that is close to $(n - d)/2$. Henceforth, let $\boldsymbol{\delta} := \mathbf{e}[\tau : \tau + d] / (\frac{m \cdot h}{4})$ be the non-zero coordinates of the correction vector \mathbf{e} scaled by $\frac{m \cdot h}{4}$.

By Equation (2), we know that the correction vector $\boldsymbol{\delta}$ must satisfy the following system of linear equations where $t := \frac{n-d}{4}$:

$$\begin{pmatrix} \binom{n-d}{\tau} & \binom{n-d}{\tau+1} & \cdots & \binom{n-d}{\tau+d} \\ \binom{n-d}{\tau-1} & \binom{n-d}{\tau} & \cdots & \binom{n-d}{\tau+d-1} \\ \vdots & \vdots & \ddots & \vdots \\ \binom{n-d}{\tau-d} & \binom{n-d}{\tau-d+1} & \cdots & \binom{n-d}{\tau} \end{pmatrix} \cdot \boldsymbol{\delta} = \begin{pmatrix} \sum_{i=0}^t \binom{n-d}{i} \\ \sum_{i=0}^{t-1} \binom{n-d}{i} \\ \vdots \\ \sum_{i=0}^{t-d} \binom{n-d}{i} \end{pmatrix}. \quad (3)$$

In Lemma 4.7, we prove that as long as $d \leq \frac{1}{8} \sqrt{\frac{h}{2 \log h}}$, and that τ is an appropriate choice close to $n/2$, then the solution $\boldsymbol{\delta}$ to the linear system in Equation (3) has a small infinity norm — specifically, $\|\boldsymbol{\delta}\|_\infty \leq 1$ — such that the resulting \mathbf{y} vector will respect the budget feasibility constraints, i.e., Equation (1). The actual proof of this bound is somewhat involved and thus deferred to Section 4.2. In particular, a key step is to bound the *smallest singular value* of the matrix in Equation (3) (henceforth denoted A) appropriately — to achieve this, we first bound A 's determinant, and then use an inequality proven by [YG97] which relates the smallest singular value and the determinant.

2.3 Finite Block Setting

2.3.1 Strict Incentive Compatibility

Feasibility for $c = 1$. The LP-based mechanism confirms any bid that offers to pay at least m . Thus, total number of confirmed bids may be unbounded. Therefore, when the block size k is finite, we cannot directly run the LP-based mechanism. We suggest the following modification to the LP-based mechanism such that it works for the finite-block setting:

MPC-assisted, LP-based mechanism with random selection

// Let k be the block size, let m be the median \mathcal{D} such that $\Pr_{x \sim \mathcal{D}}[x \geq m] = 1/2$.

- All bids offering at least m are candidates. If there are more than k candidates, randomly select k of them to confirm; else confirm all candidates. Every confirmed bid pays m .
- Let n be the length of the bid vector, let $\mathbf{y} = (y_0, y_1, \dots, y_n)$ be any feasible solution to the following linear program:

$$\forall i \in [n] : 0 \leq y_i \leq \min(i, k) \cdot m \quad (4)$$

$$\forall 0 \leq j \leq d : \sum_{i=0}^{n-d} q_i \cdot y_{i+j} = \frac{m \cdot \min(h, k)}{4} \quad (5)$$

where q_i is the probability of observing i heads if we flip $n - d$ independent fair coins.

- The total miner revenue is y_s where s is the number of *candidates*.

In comparison with the earlier LP-based mechanism, we modify the budget feasibility constraints (Equation (4)) to make sure that the total miner revenue is constrained by the actual number of confirmed bids which is now $\min(i, k)$ if the number of candidates is i . Further, we modify the expected miner revenue (Equation (5)) to be $\frac{m \cdot \min(h, k)}{4}$ which takes into account the block size k . In Section 5.1.1, we prove that as long as $c \leq d \leq \frac{1}{8} \sqrt{\frac{h}{2 \log h}}$, the above LP indeed has a feasible solution and the resulting mechanism satisfies ex post UIC, Bayesian MIC, and Bayesian SCP in $(h, *, 1, d)$ -environments.

Infeasibility for $c \geq 2$. Unfortunately, the above approach fails for $c \geq 2$. In this case, two users Alice and Bob may be in the same coalition. Alice can now help Bob simply by dropping out and not posting a bid, thus effectively increasing Bob’s chance of getting confirmed. In the event that Alice’s true value is very small and Bob’s true value is sufficiently large, this strategic action can increase the coalition’s joint utility.

Interestingly, it turns out that this is no accident. In fact, we prove that for any $h \geq 1$, $\rho \in (0, 1)$, and $d \geq c \geq 2$, no “interesting” mechanism can simultaneously achieve Bayesian UIC, MIC and SCP in (h, ρ, c, d) -environments — any such mechanism must suffer from 0 total social welfare for the users if the actual number of bids received is greater than h (see Theorem 5.2 for the formal statement). We can regard Theorem 1.6 as a generalization of [SCW23]’s Theorem 5.5: they show that any MPC-assisted mechanism that achieves Bayesian UIC, MIC, and SCP in $(*, \rho, c, *)$ -environments for $c \geq 2$ must suffer from 0 social welfare for users.

Proof roadmap. We use the following blueprint to prove Theorem 1.6. Below, consider any TFM that satisfies Bayesian UIC, MIC and Bayesian SCP in (h, ρ, c, d) -environments where $d \geq c \geq 2$.

1. First, in Lemma 5.3, using techniques inspired by Goldberg and Hartline [GH05] we prove the following: provided that there are at least h honest users (not including i and j) whose bids are sampled at random from \mathcal{D} , then a strategic user i changing its bid should not affect the utility of another user j , if user j 's bid is also sampled at random from \mathcal{D} .
2. Next, in Lemma 5.3, we prove a strategic user i dropping out should not affect another user j 's utility, assuming that at least h bids (excluding user i) sampled at random from \mathcal{D} .
3. Next, in Corollary 5.4, we show that in a world of at least h random bids (excluding user i) sampled from \mathcal{D} , user i 's expected utility when its bid is sampled randomly from \mathcal{D} depends only on i 's identity, and does not depend on the identities of the other random bids. Therefore, henceforth we can use U_i to denote this expected utility.
4. Next, in Lemma 5.5, we show that for any two identities i, j , it must be that $U_i = U_j$, otherwise, it violates the assumption that the mechanism is weakly symmetric (see definition of weak symmetry below).
5. Next, we can show that $U_i = 0$: imagine a world with K bids sampled independently from \mathcal{D} whose support is bounded. There must exist some user whose confirmation probability is upper bounded by k/K . This user's expected utility must be arbitrarily small when K is arbitrarily large. With a little more work, we can show that if the world consists of more than h bids sampled independently at random from \mathcal{D} , it must be that every user's expected utility is 0.

One technicality that arises in the full proof (see Section 5.1.2) is the usage of the weak symmetry assumption. In particular, the proof would have been much easier if we could instead assume *strong symmetry* which, unfortunately, is too stringent. In strong symmetry, we assume that any two users who bid the same amount will receive the same treatment. While it is a good approach for gaining intuition about the proof, it is too stringent since there could well be more bids offering the same value than the block size k — in this case, a non-trivial mechanism would treat them differently, i.e., confirm some while rejecting others. Our actual proof of Theorem 1.6 needs only a *weak symmetry* assumption which is a standard assumption made in prior works [CS23, SCW23], that is, if two input bid vectors \mathbf{b}, \mathbf{b}' of length n are permutations of each other, then the joint distribution of the *set* of outcomes $\{(x_i, p_i)\}_{i \in [n]}$ must be identical. This implies that if the input bid vectors are permutations of each other, then the vector of expected utilities are permutations of each other too.

2.3.2 Approximate Incentive Compatibility

Because of the limitation shown in Theorem 1.6, we relax the notion to approximate incentive compatibility, and ask if we can achieve optimal miner revenue in the finite block setting. Consider the following TFM.

MPC-assisted, diluted threshold-based Mechanism

/ Let k be the block size, let m be the median of \mathcal{D} such that $\Pr_{x \sim \mathcal{D}}[x \geq m] = 1/2$, let T be the maximum value of the distribution \mathcal{D} . */*

- Let $R := \max\left(2c\sqrt{\frac{kT}{\epsilon}}, k\right)$. All bids offering at least m are candidates. If the number of candidates $s \leq R$, randomly select $\frac{k}{R} \cdot s$ candidates to confirm; else, randomly select k candidates to confirm. Every confirmed bid pays m .

- If $s \geq \frac{h}{4}$, then the total miner revenue is $\min(\frac{h}{4} \cdot \frac{k}{R}, k) \cdot m$. Otherwise, the miners get nothing.

Intuitively, here are modifying the earlier threshold-based mechanism to 1) make it compatible with finite block size, and 2) make sure that up to c users dropping out can only minimally increase their friend’s probability of getting confirmed. In particular, resilience to drop-out is achieved by artificially diluting the probability that a user is confirmed when the number of eligible bids (i.e., offering at least m) is small. With the dilution, we guarantee that a coalition of c users cannot noticeably alter their own probability of getting confirmed, nor their friend’s probability. This implies a strategic coalition has little influence over the expected utility of all users in the coalition. Moreover, we guarantee that a strategic coalition has very little influence on the miner revenue as well: similar to the threshold-based mechanism, except with $\exp(-\Omega(h))$ probability, the miner revenue is an a-priori fixed amount, that is, $\min(\frac{h}{4} \cdot \frac{k}{R}, k) \cdot m$. Summarizing the above, we can show that the mechanism satisfies ex post UIC, Bayesian ϵ -MIC, and Bayesian ϵ -SCP in $(h, *, c, *)$ -environments, as long as $\epsilon \geq m \cdot \frac{h}{2} \cdot e^{-\frac{h}{16}}$.

Finally, for sufficiently large $h \geq \max(4k, 8c\sqrt{\frac{kT}{\epsilon}})$, the mechanism achieves $k \cdot m$ total miner revenue and $k \cdot C_{\mathcal{D}}$ user social welfare where $C_{\mathcal{D}}$ is defined in Theorem 1.1. For example, suppose we are willing to tolerate $\epsilon = 0.01T$, then we just need $h \geq \max(4k, 80c \cdot \sqrt{k})$ to achieve asymptotic optimality in miner revenue and social welfare. The full proof is deferred to Section 5.2.

2.4 Additional Results

Limit on miner revenue. In Section 6.1, we prove Theorem 1.1. Specifically, we prove that $\Theta(h)$ revenue is optimal in (h, ρ, c, d) -environments for strict incentive compatibility; and further, we generalize the bound to approximate incentive compatibility as well (see Theorem 6.1). The proof is a generalization of the techniques proposed by Shi, Chung, Wu [SCW23]. Specifically, they proved that any mechanism that satisfies Bayesian UIC, MIC, and SCP in $(*, \rho, 1, *)$ -environments must suffer from 0 miner revenue. In their proof, they argue that if we remove one bid, the miner revenue must be unaffected. In our case, because the mechanism is promised a lower bound h on the number of honest users, we can repeat this argument till there are h honest bids left, and no more. This gives rise to an $O(h)$ limit on miner revenue. The proof for the approximate incentive compatibility is also a generalization of [SCW23]’s techniques, but more technically involved since even Myerson’s lemma does not hold for approximate incentive compatibility.

Necessity of Bayesian equilibrium. As mentioned, our reasonable-world assumptions (formalized through the definition of an (h, ρ, c, d) -environment) would not have helped had we insisted on ex post notions of equilibrium (for all of UIC, MIC, and SCP). In Section 6.2, we explain why for ex post notions of incentive compatibility, even mechanisms in the (h, ρ, c, d) -environment are subject to the same miner-revenue limitations of universal mechanisms.

3 Model and Definitions

Imagine that there are n_0 users, and each user has a transaction that wants to be confirmed. For $i \in [n_0]$, let v_i be user i ’s true valuation of getting its transaction confirmed. We want to design a transaction fee mechanism (TFM) such that no individual user or miner or a coalition thereof have any incentive to deviate from honest behavior. Throughout the paper, we consider a single-parameter environment, i.e., each user’s bid is represented by a single, non-negative real number.

Chung and Shi [CS23]’s results ruled out the existence of interesting TFMs in the plain model without cryptography. First, they show a *zero miner-revenue* bound: any TFM that guarantees incentive compatibility for each individual user as well as for a miner-user coalition must suffer from zero miner revenue. The zero miner-revenue bound holds matter whether the block size is infinite or finite, and even when the miner is allowed to collude with only one user. Second, they prove a *finite-block impossibility*: assuming finite block size, then no TFM can simultaneously guarantee incentive compatibility for each individual user as well as for a miner-user coalition (even when the miner is allowed to collude with at most one user).

The subsequent work of Shi, Chung, and Wu [SCW23] considered how cryptography can help circumvent these strong impossibilities. They proposed the *MPC-assisted model*, where the rules of the TFM are enforced through a multi-party computation (MPC) protocol jointly executed among a set of miners. Unlike the plain model, the MPC-assisted model guarantees that a single miner cannot unilaterally decide which transactions to include in the block. This ties the hands of the strategic player(s), and this model indeed results in interesting mechanisms that achieve properties that would otherwise be impossible in the plain model. Below, we review the MPC-assisted model proposed by [SCW23].

3.1 MPC-Assisted Model

The definition of TFM has been given in Section 2.1. Here, we formalize the MPC-assisted model and the strategy spaces for users and miners.

Ideal-world game. Recently, blockchain projects such as Ethereum are developing “encrypted mempool” techniques which can be viewed as concrete protocols that realize an MPC-assisted model for TFM. However, for understanding the game theoretic landscape, it helps to abstract out the cryptography and think of it as a trusted *ideal functionality* (henceforth denoted \mathcal{F}_{mpc}) that always honestly implements the rules of the TFM.

With the ideal functionality \mathcal{F}_{mpc} , we can imagine the following game that captures an instance of the TFM:

1. Each user registers zero, one, or more identities with \mathcal{F}_{mpc} , and submits exactly one bid on behalf of each identity.
2. Using the vector of input bids as input, \mathcal{F}_{mpc} executes the rules of the TFM. \mathcal{F}_{mpc} now sends to all miners and users the output of the mechanism, including the set of bids that are confirmed, how much each confirmed bid pays, and how much revenue the miner gets.

We make a couple of standard assumptions:

- *Individual rationality*: each confirmed bid should pay no more than the bid itself;
- *Budget feasibility*: the miner revenue should not exceed the total payment from all confirmed bids.

Using standard techniques in cryptography, we can instantiate the ideal functionality \mathcal{F}_{mpc} using an actual cryptographic protocol among the miners and users (see Appendix D of [SCW23]). Further, in the actual instantiation, the users only need to be involved in the input phase: they only need to verifiably secret-share their input bids among all miners, and they need not be involved in the remainder of the protocol. The miners then jointly run some MPC protocols to compute securely the outcome of the auction. We can use an MPC protocol that retains security even when

all but one miner are corrupt [GMW87]. Such protocols achieve a security notion called “security with abort”, i.e., an adversary controlling a majority coalition can cause the protocol to abort without producing any outcome. Conceptually, one can imagine that in the ideal-world protocol where parties interact with \mathcal{F}_{mpc} , the adversary is allowed to send \perp to \mathcal{F}_{mpc} , in which case \mathcal{F}_{mpc} will abort and output \perp to everyone. However, no strategic coalition should have an incentive to cause the protocol to abort — in this case, no block will be mined and the coalition has a utility of 0. Thus, without loss of generality, we need not explicitly capture aborting as a possible strategy in our ideal-world game mentioned above.

Strategy space and utility. An honest user will always register a single identity and submit only one bid reflecting its true value. A strategic user or miner (possibly colluding with others) can adopt the following strategies or a combination thereof:

- *Bid untruthfully*: a strategic user can misreport its value;
- *Inject fake bids*: a strategic user or miner can *inject fake bids* by registering fake identities;
- *Drop out*: a strategic user can also *drop out* by not registering its real identity.

In the real-world cryptographic instantiation, strategic miners can also deviate from the honest MPC protocol. However, as mentioned, the MPC protocol retains security (i.e., can be simulated by the ideal-world game) as long as at least one miner is honest. Therefore, we need not explicitly capture this deviation in the ideal-world game. Finally, strategic miners can cause the MPC protocol to abort without producing output, and as mentioned, this deviation never makes sense since it results in a utility of 0; thus we also need not explicitly capture it in the ideal-world game.

Let v_i denote user i ’s true value. If user i ’s transaction is confirmed and its payment is p_i , then its utility is defined as $v_i - p_i$. The miner’s utility is simply its revenue. The joint utility of a coalition \mathcal{C} is defined as the sum of the utilities of all members in \mathcal{C} .

3.2 Defining Incentive Compatibility

In the plain model without cryptography, users submit their bids in the clear over a broadcast channel, and a strategic coalition can decide its strategy after observing the remaining honest users’ bids. By contrast, in the MPC-assisted model, bids are submitted to the ideal functionality \mathcal{F}_{mpc} (in the actual cryptographic instantiation, the users verifiably secret-share their bids among the miners). This means that the strategic coalition must now submit its bids without having observed other users’ bids. Therefore, in the MPC model, it makes sense to consider a *Bayesian* notion of equilibrium rather than an *ex post* notion. In an *ex post* setting, we require that a strategic individual or coalition’s best response is to act honestly even after having observed others’ actions. In a *Bayesian* setting, we assume that every honest user’s bid is sampled independently from some distribution \mathcal{D} , and we require that acting honestly maximizes the strategic individual or coalition’s expected utility where the expectation is taken over not just the random coins of the TFM itself, but also over the randomness in sampling the honest users’ bids.

Notations. Henceforth, we use the notation \mathbf{b} to denote a bid vector. Since we allow strategic players to inject fake bids or drop out, the length of \mathbf{b} need not be the same as the number of users n_0 . We use the notation \mathcal{C} to denote a coalition, and we use $\mathbf{b}_{-\mathcal{C}}$ to denote the bid vector belonging to honest users outside \mathcal{C} . We use the notation $\mathcal{D}_{-\mathcal{C}}$ to denote the joint distribution of $\mathbf{b}_{-\mathcal{C}}$, that is, $\mathcal{D}_{-\mathcal{C}} = \mathcal{D}^{h_0}$ where h_0 denotes the number of honest users outside \mathcal{C} . Similarly, if i is an individual

strategic user, then the notation \mathbf{b}_{-i} denotes the bid vector belonging the remaining honest users in $[n_0] \setminus \{i\}$. We use the notation \mathcal{D}_{-i} to denote the joint distribution of the honest bid vector \mathbf{b}_{-i} .

For generality, we define *approximate* incentive compatibility parameterized by an additive slack ϵ . The case of *strict* incentive compatibility can be viewed as the special case when $\epsilon = 0$.

3.2.1 Bayesian Incentive Compatibility

Definition 3.1 (Bayesian incentive compatibility). We say that an MPC-assisted TFM satisfies Bayesian ϵ -incentive compatibility for a coalition or individual \mathcal{C} , iff for any $\mathbf{v}_{\mathcal{C}}$ denoting the true values of users in \mathcal{C} , sample $\mathbf{b}_{-\mathcal{C}} \sim \mathcal{D}_{-\mathcal{C}}$, then, no strategy can increase \mathcal{C} 's expected utility by more than ϵ in comparison with honest behavior, where the expectation is taken over randomness of the honest users' bids $\mathbf{b}_{-\mathcal{C}}$, as well as random coins consumed by the TFM. Specifically, we define the following notions depending on who is the strategic individual or coalition:

- *User incentive compatibility (UIC)*. We say that an MPC-assisted TFM satisfies Bayesian ϵ -UIC in some environment \mathcal{E} , iff for any n , for any user $i \in [n]$, for any true value $v_i \in \mathbb{R}^{\geq 0}$ of user i , for any strategic bid vector \mathbf{b}_i from user i which could be empty or consist of multiple bids, the following holds as long as the conditions required by the environment \mathcal{E} are respected:

$$\mathbf{E}_{\mathbf{b}_{-i} \sim \mathcal{D}_{-i}} [\text{util}^i(\mathbf{b}_{-i}, v_i)] \geq \mathbf{E}_{\mathbf{b}_{-i} \sim \mathcal{D}_{-i}} [\text{util}^i(\mathbf{b}_{-i}, \mathbf{b}_i)] - \epsilon$$

where $\text{util}^i(\mathbf{b})$ denotes the expected utility (taken over the random coins of the TFM) of user i when the bid vector is \mathbf{b} .

- *Miner incentive compatibility (MIC)*. We say that an MPC-assisted TFM satisfies Bayesian ϵ -MIC in some environment \mathcal{E} , iff for any miner coalition \mathcal{C} , for any strategic bid vector \mathbf{b}' injected by the miner, the following holds as long as the conditions required by the environment \mathcal{E} are respected:

$$\mathbf{E}_{\mathbf{b}_{-\mathcal{C}} \sim \mathcal{D}_{-\mathcal{C}}} [\text{util}^{\mathcal{C}}(\mathbf{b}_{-\mathcal{C}})] \geq \mathbf{E}_{\mathbf{b}_{-\mathcal{C}} \sim \mathcal{D}_{-\mathcal{C}}} [\text{util}^{\mathcal{C}}(\mathbf{b}_{-\mathcal{C}}, \mathbf{b}')] - \epsilon$$

where $\text{util}^{\mathcal{C}}(\mathbf{b})$ denotes the expected utility (taken over the random coins of the TFM) of the coalition \mathcal{C} when the input bid vector is \mathbf{b} .

- *Side-contract-proofness (SCP)*. We say that an MPC-assisted TFM satisfies Bayesian ϵ -SCP in some environment \mathcal{E} , iff for any miner-user coalition, for any true value vector $\mathbf{v}_{\mathcal{C}}$ of users in \mathcal{C} , for any strategic bid vector $\mathbf{b}_{\mathcal{C}}$ of the coalition (whose length may not be equal to the number of users in \mathcal{C}), the following holds as long as the requirements of the environment \mathcal{E} are respected:

$$\mathbf{E}_{\mathbf{b}_{-\mathcal{C}} \sim \mathcal{D}_{-\mathcal{C}}} [\text{util}^{\mathcal{C}}(\mathbf{b}_{-\mathcal{C}}, \mathbf{v}_{\mathcal{C}})] \geq \mathbf{E}_{\mathbf{b}_{-\mathcal{C}} \sim \mathcal{D}_{-\mathcal{C}}} [\text{util}^{\mathcal{C}}(\mathbf{b}_{-\mathcal{C}}, \mathbf{b}_{\mathcal{C}})] - \epsilon$$

Henceforth, if a mechanism satisfies Bayesian ϵ -UIC for $\epsilon = 0$ (i.e., the *strict* incentive compatibility case), we often omit writing the ϵ , and simply say that the mechanism satisfies Bayesian UIC. The terms ‘‘Bayesian MIC’’, and ‘‘Bayesian SCP’’ are similarly defined.

Notice that we only require honest users' true values are i.i.d. sampled, while the strategic players' true values can be arbitrary.

3.2.2 Ex Post Incentive Compatibility

Definition 3.2 (Ex post incentive compatibility). We say that a TFM satisfies ex post ϵ -UIC, ϵ -MIC, and ϵ -SCP respectively, for a coalition or individual \mathcal{C} , iff the following conditions hold, respectively:

- *User incentive compatibility (UIC)*. We say that an MPC-assisted TFM satisfies ex post ϵ -UIC in some environment \mathcal{E} , iff for any n , for any user $i \in [n]$, for any bid vector \mathbf{b}_{-i} denoting the bids of everyone else besides i , for any true value $v_i \in \mathbb{R}^{\geq 0}$ of user i , for any strategic bid vector \mathbf{b}_i from user i which could be empty or consist of multiple bids, the following holds as long as the conditions required by the environment \mathcal{E} are respected:

$$\text{util}^i(\mathbf{b}_{-i}, v_i) \geq \text{util}^i(\mathbf{b}_{-i}, \mathbf{b}_i) - \epsilon$$

where $\text{util}^i(\mathbf{b})$ denotes the expected utility (taken over the random coins of the TFM) of user i when the bid vector is \mathbf{b} .

- *Miner incentive compatibility (MIC)*. We say that an MPC-assisted TFM satisfies ex post ϵ -MIC in some environment \mathcal{E} , iff for any miner coalition \mathcal{C} , for any bid vector $\mathbf{b}_{-\mathcal{C}}$, for any strategic bid vector \mathbf{b}' injected by the miner, the following holds as long as the conditions required by the environment \mathcal{E} are respected:

$$\text{util}^{\mathcal{C}}(\mathbf{b}_{-\mathcal{C}}) \geq \text{util}^{\mathcal{C}}(\mathbf{b}_{-\mathcal{C}}, \mathbf{b}') - \epsilon$$

where $\text{util}^{\mathcal{C}}(\mathbf{b})$ denotes the expected utility (taken over the random coins of the TFM) of the coalition \mathcal{C} when the input bid vector is \mathbf{b} .

- *Side-contract-proofness (SCP)*. We say that an MPC-assisted TFM satisfies ex post ϵ -SCP in some environment \mathcal{E} , iff for any miner-user coalition, for any bid vector $\mathbf{b}_{-\mathcal{C}}$ submitted by non-coalition-members, for any true value vector $\mathbf{v}_{\mathcal{C}}$ of users in \mathcal{C} , for any strategic bid vector $\mathbf{b}_{\mathcal{C}}$ of the coalition (whose length may not be equal to the number of users in \mathcal{C}), the following holds as long as the requirements of the environment \mathcal{E} are respected:

$$\text{util}^{\mathcal{C}}(\mathbf{b}_{-\mathcal{C}}, \mathbf{v}_{\mathcal{C}}) \geq \text{util}^{\mathcal{C}}(\mathbf{b}_{-\mathcal{C}}, \mathbf{b}_{\mathcal{C}}) - \epsilon$$

Henceforth, if a mechanism satisfies ex post ϵ -UIC for $\epsilon = 0$ (i.e., the *strict* incentive compatibility case), we often omit writing the ϵ , and simply say that the mechanism satisfies ex post UIC. The terms “ex post MIC”, “ex post SCP” are similarly defined.

Recall that in the game representing the plain model, strategic players can choose their actions *after* having observed the bids submitted by honest users. This gives rise to the following fact which essentially says it does not make sense to consider Bayesian notions of equilibrium in the plain model.

Fact 3.3. *Any plain-model TFM that satisfies Bayesian ϵ -UIC (or Bayesian ϵ -MIC, Bayesian ϵ -SCP resp.) in some environment \mathcal{E} must also satisfy ex post ϵ -UIC (or ex post ϵ -MIC, ex post ϵ -SCP resp.).*

4 Feasibility for Infinite Block Size

4.1 MPC-Assisted, Threshold-Based Mechanism

We assume that honest users’ bids are drawn i.i.d. from some distribution \mathcal{D} with the median m such that $\Pr_{x \sim \mathcal{D}}[x \geq m] = 1/2$ (see Remark 2.1). For convenience, we repeat the MPC-assisted, threshold-based mechanism, which has been introduced in Section 2.2.

MPC-assisted, threshold-based mechanism

Parameters: lower bound h on the number of honest users, the distribution median m .

Mechanism:

- *Confirmation rule.* Given a bid vector $\mathbf{b} = (b_1, \dots, b_\ell)$, for each bid b_i , confirm b_i if $b_i \geq m$.
- *Payment rule.* Each confirmed bid pays m .
- *Miner revenue rule.* Let s be the number of confirmed bids. If $s \geq \frac{h}{4}$, miner gets $\frac{h}{4} \cdot m$. Otherwise, the miner gets nothing.

Theorem 4.1. *Fix any $h \geq 1$. The MPC-assisted, threshold-based mechanism satisfies ex post UIC, Bayesian ϵ -MIC and Bayesian ϵ -SCP in an $(h, *, *, *)$ -environment, where $\epsilon = \frac{h}{4} \cdot m \cdot e^{-\frac{h}{16}}$.*

Proof. First, UIC follows from the same reasoning as in Theorem A.1. We will focus on MIC and SCP in the rest of the proof.

ϵ -MIC. Recall that the only strategy that a strategic miner can apply is injecting some fake bids. Because injecting fake bids smaller than m does not influence the colluding miners' utility, we only consider injecting bids at least m . Let X denote the random variable representing the number of honest bids at least m . The only situation where the colluding miners can increase their expected gain by injecting fake bids is when $X < \frac{h}{4}$. By the following Chernoff Bound,

Lemma 4.2 (Chernoff bound, Corollary A.1.14 [AS16]). *Let X_1, \dots, X_n be independent Bernoulli random variables. Let $\mu = \mathbf{E}[\sum_{i=1}^n X_i]$. Then, for any $\epsilon \in (0, 1)$, it holds that*

$$\Pr \left[\sum_{i=1}^n X_i \leq (1 - \epsilon)\mu \right] \leq e^{-\epsilon^2 \mu / 2}.$$

We have

$$\Pr \left[X < \frac{h}{4} \right] \leq e^{-\frac{h}{16}}.$$

Therefore, the colluding miners can gain at most $\frac{h}{4} \cdot m \cdot e^{-\frac{h}{16}}$ more expected revenue by injecting fake bids.

ϵ -SCP. Since the confirmation and the payment of each bid are independent of other bids, and the mechanism is strict UIC, the coalition cannot increase colluding users' utilities. Therefore, by deviating from the mechanism, the coalition can only try to increase the expected total miner revenue. By a similar argument as MIC, the coalition can only increase the expected total miner revenue when $X < \frac{h}{4}$, which happens with a probability no more than $e^{-\frac{h}{16}}$. It follows that no matter how the coalition deviates, the expected miner's revenue can increase by at most $\rho \cdot \frac{h}{4} \cdot m \cdot e^{-\frac{h}{16}}$. □

4.2 Analysis of the LP-Based Mechanism

4.2.1 Preliminaries: Linear Algebra Tools

We first introduce some linear algebra tools needed for analyzing the LP-based mechanism.

Throughout this section, all our indexing for vectors and matrices starts from 0. Given a vector $\mathbf{b} = (b_0, b_1, \dots, b_n)$ and two integers i, j such that $i \leq j$, we define $\mathbf{b}[i : j]$ to be the subvector

(b_i, \dots, b_j) . We use $A = (a_{ij}) \in \mathbb{R}^{n,m}$ to denote a matrix in which the entry of the i -th row and j -th column is a_{ij} . Let A^T denote the transpose of A , and A^{-1} denote the inverse of A if A is non-singular.

Norm. Define the *infinity-norm* $\|\mathbf{b}\|_\infty$ of a vector \mathbf{b} to be $\|\mathbf{b}\|_\infty = \max\{|b_i| : 0 \leq i \leq n\}$. For a square $n \times n$ matrix $A = (a_{ij})$, define the following matrix norms:

- *Infinity norm:* $\|A\|_\infty = \sup_{\|x\|_\infty=1} \|Ax\|_\infty = \max_i \sum_{j=1}^n |a_{ij}|$.
- *ℓ_2 -norm:* $\|A\|_2 = \sup_{\|x\|_2=1} \|Ax\|_2$.
- *Frobenius norm:* $\|A\|_F = \left(\sum_{i,j=0}^{n-1} a_{ij}^2 \right)^{1/2}$.

It is easy to check that $\|A\|_\infty \leq \|A\|_2$, and that $\|Ax\|_\infty \leq \|A\|_\infty \|x\|_\infty$.

Singular value. For a square $n \times n$ matrix A , the singular values are the square roots of the eigenvalues of $A^T A$.

Fact 4.3. Let $A \in \mathbb{R}^{n \times n}$ be non-singular. Let $\lambda_1 \geq \dots \geq \lambda_n$ be the singular values of A . Then $\|A^{-1}\|_2 = \frac{1}{\lambda_n}$.

Lemma 4.4 (Yu and Gu [YG97]). Let $A \in \mathbb{R}^{n \times n}$ be non-singular and λ be the smallest singular value of A . Then

$$\lambda \geq |\det(A)| \cdot \left(\frac{n-1}{\|A\|_F^2} \right)^{(n-1)/2} > 0.$$

Determinant. The determinant of a matrix $A = (a_{ij}) \in \mathbb{R}^{n \times n}$ is $\det(A) = \sum_{\sigma \in S_n} \text{sgn}(\sigma) \prod_{i=1}^n a_{i,\sigma_i}$, where S_n is the set of all permutations σ over the set $\{0, \dots, n-1\}$. For each permutation $\sigma \in S_n$, let σ_i denote the value of the i -th position after reordering by σ . The signature $\text{sgn}(\sigma)$ of a permutation σ is $+1$ if the permutation can be obtained by an even number of swaps between two entries and -1 otherwise.

4.2.2 Proofs for the LP-Based Mechanism

We now prove that the MPC-assisted LP-based mechanism satisfies strict incentive compatibility in an $(h, *, c, d)$ -environment. Suppose that honest users' values are sampled i.i.d. from some distribution \mathcal{D} . Recall that m denotes the median of the bid distribution \mathcal{D} , and $C_{\mathcal{D}} = \mathbf{E}_{x \sim \mathcal{D}}[x - m | x \geq m]$ is another constant related to the distribution \mathcal{D} . Without loss of generality, we assume $\Pr[x \geq m] = \frac{1}{2}$ (see Remark 2.1).

Theorem 4.5 (Theorem 1.3 restated). *Suppose that the block size is infinite. Fix any⁷ $h \geq 2$, and any $d \leq \frac{1}{8} \sqrt{\frac{h}{2 \log h}}$, the MPC-assisted, LP-based mechanism guarantees ex post UIC, Bayesian MIC, and Bayesian SCP in an $(h, *, *, d)$ -environment, and meanwhile, the mechanism achieves $\Theta(h \cdot m)$ expected miner revenue, and at least $\Theta(\tilde{h} \cdot C_{\mathcal{D}})$ expected social welfare for the users where $\tilde{h} \geq h$ is the actual number of honest users that show up.*

⁷For the special case $h = 1$, we can just use the parity-based mechanism of Section 2.2.

We prove Theorem 4.5 in two steps. First, we show that if the linear program defined in Equations (1) and (2) has a feasible solution, then the resulting mechanism satisfies incentive compatibility, as formally stated below:

Lemma 4.6. *When the linear program defined in Equations (1) and (2) has a feasible solution, the LP-based mechanism satisfies ex post UIC, Bayesian MIC, and Bayesian SCP in an $(h, *, *, d)$ -environment. Moreover, the expected miner revenue is $\frac{h \cdot m}{4}$, and the user social welfare is $\Theta(\tilde{h} \cdot C_{\mathcal{D}})$.*

Proof. First, it is easy to see that the expected total miner revenue is $\frac{h \cdot m}{4}$, as guaranteed by the linear program Equations (1) and (2). Moreover, since the expected utility of a user with true value v is $v - m$ if $v \geq m$, the expected user social welfare is at least

$$\sum_{i \in H} \mathbf{E}_{v_i \sim \mathcal{D}} [v_i - m \mid v_i > m] = \tilde{h} \cdot \mathbf{E}_{x \sim \mathcal{D}} [x - m \mid x \geq m],$$

where H is the set of all honest users.

Next, we prove that the mechanism is strict incentive compatible if the linear program has a solution. UIC is easy to see. Next, we only prove SCP since MIC follows from the same reasoning.

SCP. Since the confirmation and the payment of each bid are independent of other bids, and the mechanism is strict UIC, the coalition cannot increase colluding users' expected utilities. Therefore, we only need to show that the coalition cannot increase the expected total miner revenue by deviating from the mechanism. Intuitively, the linear program Equations (1) and (2) ensures that for arbitrary d bids, the total miner revenue taking an expectation over the remaining $n - d$ bids always remains $\frac{h \cdot m}{4}$.

Formally, let \tilde{h} denote the number of *real honest bids* and \mathbf{b}_{-c} denote the random variable of honest users' bids. Then $\tilde{h} \geq n - d = \gamma$. For any bid \mathbf{b}_c controlled by the coalition, the expected total miner revenue is

$$\mathbf{E}_{\mathbf{b}_{-c} \sim \mathcal{D}^{\tilde{h}}} [\mu(\mathbf{b}_{-c}, \mathbf{b}_c)] = \int_{\mathbf{t} \sim \mathcal{D}^{\tilde{h}-\gamma}} \mathbf{E}_{\mathbf{b} \sim \mathcal{D}^{\gamma}} [\mu(\mathbf{b}, \mathbf{t}, \mathbf{b}_c)] f(\mathbf{t}) dt, \quad (6)$$

where $f(\cdot)$ is the p.d.f. for $\mathcal{D}^{\tilde{h}-\gamma}$. For any fixed $(\mathbf{t}, \mathbf{b}_c)$, let I denote the number of bids that are larger than or equal to m in $(\mathbf{t}, \mathbf{b}_c)$. Since the probability of an honest bid being at least m is exactly $\frac{1}{2}$,

$$\mathbf{E}_{\mathbf{b} \sim \mathcal{D}^{\gamma}} [\mu(\mathbf{b}, \mathbf{t}, \mathbf{b}_c)] = \sum_{i=0}^{\gamma} \frac{1}{2^{\gamma}} \binom{\gamma}{i} y_{i+I},$$

which is exactly $\frac{h \cdot m}{4}$ as guaranteed by Equation (2). Substituting back into (6), for any bid \mathbf{b}_c , we have that

$$\mathbf{E}_{\mathbf{b}_{-c} \sim \mathcal{D}^{\tilde{h}}} [\mu(\mathbf{b}_{-c}, \mathbf{b}_c)] = \int_{\mathbf{t} \sim \mathcal{D}^{\tilde{h}-\gamma}} \frac{h \cdot m}{4} \cdot f(\mathbf{t}) dt = \frac{h \cdot m}{4}.$$

Therefore, for any d bids controlled by the coalition, the expected miner revenue remains $\frac{h \cdot m}{4}$. \square

In the main body, we focus on proving the more challenging step, that is, as long as $d \leq \frac{1}{8} \sqrt{\frac{h}{2 \log h}}$, the linear program indeed has a feasible solution, formally stated below.

Lemma 4.7. For $h \geq 2$ and $d \leq \frac{1}{8} \sqrt{\frac{h}{2 \log h}}$, the linear program specified by Equations (1) and (2) is guaranteed to have a feasible solution.

Proof. We will give a constructive solution to the linear program Equations (1) and (2). Let $\gamma := n - d$ denote the number of bids that are sampled randomly from \mathcal{D} . Let $t = \lfloor \frac{\gamma}{4} \rfloor$, and $\bar{\mu}$ be our target expected miner revenue $\frac{m \cdot h}{4}$. We start from an “approximate” solution $\hat{\mathbf{y}} = (\hat{y}_0, \dots, \hat{y}_n) \in \mathbb{R}^{n+1}$ such that $\hat{y}_i = 0$ for any $i \leq t$, and $\hat{y}_i = \bar{\mu}$ for any $i > t$. Our goal is to find a correction $\mathbf{e} = (e_0, \dots, e_n) \in \mathbb{R}^{n+1}$ that is zero everywhere except for the indices $i \in [z + d, z + 2d]$ for some $z \geq \frac{\gamma}{2}$ such that $\hat{\mathbf{y}} + \mathbf{e}$ is a feasible solution to the linear program Equations (1) and (2). Henceforth, let $\boldsymbol{\delta} := \mathbf{e}[z + d, z + 2d] / \bar{\mu}$ be the non-zero coordinates of the correction, scaled by $\bar{\mu}$. Then $\boldsymbol{\delta}$ must satisfy the linear system $A(z)\boldsymbol{\delta} = \boldsymbol{\Delta}$, where $A(z)$ and $\boldsymbol{\Delta}$ are defined as follows:

$$A(z) = \begin{pmatrix} \binom{\gamma}{z+d} & \binom{\gamma}{z+d+1} & \cdots & \binom{\gamma}{z+2d} \\ \binom{\gamma}{z+d-1} & \binom{\gamma}{z+d} & \cdots & \binom{\gamma}{z+2d-1} \\ \vdots & \vdots & \ddots & \vdots \\ \binom{\gamma}{z} & \binom{\gamma}{z+1} & \cdots & \binom{\gamma}{z+d} \end{pmatrix}, \quad \boldsymbol{\Delta} = \begin{pmatrix} \sum_{i=0}^t \binom{\gamma}{i} \\ \sum_{i=0}^{t-1} \binom{\gamma}{i} \\ \vdots \\ \sum_{i=0}^{t-d} \binom{\gamma}{i} \end{pmatrix}.$$

If there exists a $z^* \in [\lceil \frac{\gamma}{2} \rceil, \lceil \frac{\gamma}{2} \rceil + 2d^2]$ such that this linear system $A(z^*)\boldsymbol{\delta} = \boldsymbol{\Delta}$ has a solution $\boldsymbol{\delta}$, then choosing \mathbf{e} such that $\mathbf{e}[z^* + d : z^* + 2d] = \bar{\mu} \cdot \boldsymbol{\delta}$ gives a solution $\hat{\mathbf{y}} + \mathbf{e}$ that satisfies Equation (2).

Claim 4.8. There exists a $z^* \in [\lceil \frac{\gamma}{2} \rceil, \lceil \frac{\gamma}{2} \rceil + 2d^2]$ such that the matrix $A(z^*)$ is non-singular, and

$$\|A(z^*)^{-1}\|_{\infty} \leq \frac{(z^* + 2d)^{2d(d+1)}}{\binom{\gamma}{z^*}} \cdot \left(\frac{d+1}{\sqrt{d}}\right)^d. \quad (7)$$

When choosing this z^* , we have a unique solution $\boldsymbol{\delta} = A(z^*)^{-1}\boldsymbol{\Delta}$. Moreover, under the given parameter range, the solution $\boldsymbol{\delta}$ has bounded infinity norm:

Claim 4.9. For $h \geq 2$ and $d \leq \frac{1}{8} \sqrt{\frac{h}{2 \log h}}$, we have $\|\boldsymbol{\delta}\|_{\infty} \leq 1$.

For now, we assume that Claim 4.8 and Claim 4.9 are true, and we show how they lead to Lemma 4.7. The proofs of the two claims appear right afterward. To prove Lemma 4.7, it suffices to show that $\hat{\mathbf{y}} + \mathbf{e}$ indeed satisfies the budget feasibility specified by Equation (1). Since for all $i \notin [z^* + d, z^* + 2d]$, we have $\hat{y}_i + e_i = \hat{y}_i \leq i \cdot m$, so we only need to show that for the correction position $z^* + d, \dots, z^* + 2d$, the budget feasibility is satisfied. Substituting $\|\boldsymbol{\delta}\|_{\infty} \leq 1$, for each $i \in [z^* + d, z^* + 2d]$, we have $|e_i| \leq \bar{\mu}$. This implies that $\hat{y}_i + e_i \geq \bar{\mu} - \bar{\mu} = 0$. Moreover,

$$\hat{y}_i + e_i \leq 2\bar{\mu} \leq \frac{\gamma}{2} \cdot m \leq i \cdot m.$$

Lemma 4.7 thus follows. □

Proof of Claim 4.8. We separate the proof in two parts: we first show that there exists a $z^* \in [\lceil \frac{\gamma}{2} \rceil, \lceil \frac{\gamma}{2} \rceil + 2d^2]$ such that $A(z^*)$ is non-singular; then we show that the infinity norm of the inverse of $A(z^*)$ satisfies Equation (7).

Non-singularity. We show that there exists z^* in the given range such that $\det(A(z^*)) \neq 0$. Define

$$B(z) = \frac{A(z)}{\binom{\gamma}{z}} \cdot \prod_{i=1}^{2d} (z+i).$$

Since

$$\begin{aligned} \frac{\binom{\gamma}{z+j}}{\binom{\gamma}{z}} \cdot \prod_{i=1}^{2d} (z+i) &= \frac{(\gamma-z-j+1) \dots (\gamma-z)}{(z+1) \dots (z+j)} \cdot \prod_{i=1}^{2d} (z+i) \\ &= \prod_{i=1}^j (\gamma-z-j+i) \cdot \prod_{i=j+1}^{2d} (z+i), \end{aligned}$$

$B(z)$ is equal to the following matrix:

$$\begin{pmatrix} \prod_{i=1}^d (\gamma-z-d+i) \prod_{i=d+1}^{2d} (z+i) & \prod_{i=1}^{d+1} (\gamma-z-d-1+i) \prod_{i=d+2}^{2d} (z+i) & \dots & \prod_{i=1}^{2d} (\gamma-z-2d+i) \\ \prod_{i=1}^{d-1} (\gamma-z-d+1+i) \prod_{i=d}^{2d} (z+i) & \prod_{i=1}^d (\gamma-z-d+i) \prod_{i=d+1}^{2d} (z+i) & \dots & \prod_{i=1}^{2d-1} (\gamma-z-d+1+i)(z+2d) \\ \vdots & \vdots & \ddots & \vdots \\ \prod_{i=1}^{2d} (z+i) & (\gamma-z) \prod_{i=2}^{2d} (z+i) & \dots & \prod_{i=1}^d (\gamma-z-d+i) \prod_{i=d+1}^{2d} (z+i) \end{pmatrix}$$

It is sufficient to show that there exists a $z^* \in [\lceil \frac{\gamma}{2} \rceil, \lceil \frac{\gamma}{2} \rceil + 2d^2]$ such that $\det(B(z^*)) \neq 0$. To show this, note that the determinant of $B(z)$ is a polynomial $q(z)$ of z with a degree at most $2d^2$. As long as $q(z)$ is not a zero polynomial, $q(z)$ has at most $2d^2$ roots. That means, there must exist a $z^* \in [\lceil \frac{\gamma}{2} \rceil, \lceil \frac{\gamma}{2} \rceil + 2d^2]$ such that $q(z^*) \neq 0$. The non-singularity of $A(z^*)$ thus follows.

Hence, it suffices to show that $q(z)$ is not a zero polynomial. Indeed, when $z = \gamma - d$, the matrix $B(z)$ becomes the following lower triangle matrix, which has a positive determinant.

$$\begin{pmatrix} \prod_{i=1}^c i \cdot \prod_{i=c+1}^{2c} (z+i) & 0 & \dots & 0 \\ \prod_{i=1}^{c-1} (i+1) \cdot \prod_{i=c}^{2c} (z+i) & \prod_{i=1}^c i \cdot \prod_{i=c+1}^{2c} (z+i) & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \prod_{i=1}^{2c} (z+i) & c \cdot \prod_{i=2}^{2c} (z+i) & \dots & \prod_{i=1}^c i \cdot \prod_{i=c+1}^{2c} (z+i) \end{pmatrix}$$

This implies that $q(z)$ is not a zero polynomial.

Infinity norm. For simplicity, we use $A := A(z^*)$ in this part. By Fact 4.3, $\|A^{-1}\|_2 = \frac{1}{\lambda}$, where λ is the smallest singular value of A . By Lemma 4.4, the smallest singular value λ satisfies

$$\lambda \geq |\det(A)| \cdot \left(\frac{d}{\|A\|_F^2} \right)^{\frac{d}{2}}.$$

By the definition of Frobenius norm and the fact that the largest term in A is $\binom{\gamma}{z^*}$,

$$\|A\|_F^2 = \sum_{i=0}^c \sum_{j=0}^d a_{ij}^2 \leq (d+1)^2 \cdot \left(\binom{\gamma}{z^*} \right)^2.$$

We only need to bound the determinant of A . Let $A' = (a'_{i,j})_{(d+1) \times (d+1)}$ where $a'_{i,j} = \frac{a_{i,j}}{\binom{\gamma}{z^*}}$. Then we have that $|\det(A)| = \binom{\gamma}{z^*}^{(d+1)} \cdot |\det(A')|$, where

$$A' = \begin{pmatrix} \frac{(\gamma-z^*-d+1)\dots(\gamma-z^*)}{(z^*+1)\dots(z^*+d)} & \frac{(\gamma-z^*-d)\dots(\gamma-z^*)}{(z^*+1)\dots(z^*+d+1)} & \cdots & \frac{(\gamma-z^*-2d+1)\dots(\gamma-z^*)}{(z^*+1)\dots(z^*+2d)} \\ \frac{(\gamma-z^*-d+2)\dots(\gamma-z^*)}{(z^*+1)\dots(z^*+d-1)} & \frac{(\gamma-z^*-d+1)\dots(\gamma-z^*)}{(z^*+1)\dots(z^*+d)} & \cdots & \frac{(\gamma-z^*-2d+2)\dots(\gamma-z^*)}{(z^*+1)\dots(z^*+2d-1)} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & \frac{\gamma-z^*}{z^*+1} & \cdots & \frac{(\gamma-z^*-d+1)\dots(\gamma-z^*)}{(z^*+1)\dots(z^*+d)} \end{pmatrix}$$

By the definition of determinant, we have

$$\det(A') = \sum_{\sigma \in S_{d+1}} \text{sgn}(\sigma) \prod_{i=1}^{d+1} a'_{i,\sigma_i}.$$

For each $\sigma \in S_{d+1}$, let q_σ denote the product $\prod_{i=1}^{d+1} a'_{i,\sigma_i}$. Since all the entries in A' are rational numbers, for each σ , the product q_σ is also a rational number. Note that the denominator of each entry is a factor of $\prod_{i=1}^{2d} (z^* + i)$, we can write each $q_\sigma = \frac{p_\sigma}{\prod_{i=1}^{2d} (z^* + i)^{(d+1)}}$ for an integer p_σ . Thus,

$$\begin{aligned} |\det(A')| &= \left| \sum_{\sigma \in S_{d+1}} \text{sgn}(\sigma) \prod_{i=1}^{d+1} a'_{i,\sigma_i} \right| \\ &= \left| \sum_{\sigma \in S_{d+1}} \text{sgn}(\sigma) \frac{p_\sigma}{\prod_{i=1}^{2d} (z^* + i)^{(d+1)}} \right| \\ &= \frac{|\sum_{\sigma \in S_{d+1}} \text{sgn}(\sigma) p_\sigma|}{\prod_{i=1}^{2d} (z^* + i)^{(d+1)}} \geq \frac{1}{\prod_{i=1}^d (z^* + i)^{(d+1)}}, \end{aligned}$$

where the last step follows from the fact that A' is non-singular; thus the absolute value of the nominator is at least 1. Therefore,

$$\begin{aligned} \lambda &\geq |\det(A)| \cdot \left(\frac{d}{\|A\|_F^2} \right)^{\frac{d}{2}} \\ &\geq \binom{\gamma}{z^*}^{(d+1)} \cdot \frac{1}{\prod_{i=1}^{2d} (z^* + i)^{(d+1)}} \cdot \left(\frac{d}{(d+1)^2 \binom{\gamma}{z^*}^2} \right)^{\frac{d}{2}} \\ &\geq \binom{\gamma}{z^*} \cdot \frac{1}{(z^* + 2d)^{2d(d+1)}} \cdot \left(\frac{\sqrt{d}}{d+1} \right)^d. \end{aligned}$$

The claim thus follows from the fact that $\|A^{-1}\|_\infty \leq \|A^{-1}\|_2 = \frac{1}{\lambda}$. \square

Proof of Claim 4.9. Since $A(z^*)$ is non-singular, we have $\delta = A(z^*)^{-1} \Delta$. By properties of matrix norms, we have that

$$\|\delta\|_\infty \leq \|A(z^*)^{-1}\|_\infty \cdot \|\Delta\|_\infty. \quad (8)$$

By Claim 4.8 and note that $\|\Delta\|_\infty \leq t \cdot \binom{\gamma}{t}$, we have

$$\|\delta\|_\infty \leq \|A(z^*)^{-1}\|_\infty \cdot \|\Delta\|_\infty \leq \frac{(z^* + 2d)^{2d(d+1)}}{\binom{\gamma}{z^*}} \cdot \left(\frac{d+1}{\sqrt{d}}\right)^d \cdot t \cdot \binom{\gamma}{t}.$$

Because $z^* + 2d \leq \gamma$, $\frac{d+1}{\sqrt{d}} \leq \gamma$ and $t \leq \gamma$, we have

$$\|\delta\|_\infty \leq \frac{(z^* + 2d)^{2d(d+1)}}{\binom{\gamma}{z^*}} \cdot \left(\frac{d+1}{\sqrt{d}}\right)^d \cdot t \cdot \binom{\gamma}{t} \leq \gamma^{2d^2+3d+1} \cdot \frac{\binom{\gamma}{t}}{\binom{\gamma}{z^*}}. \quad (9)$$

By the assumption that $d \leq \frac{1}{8}\sqrt{\frac{h}{2\log h}}$, we have

$$(2d^2 + 3d + 1) \cdot \log(h - d) + \frac{\log \frac{6}{5}}{4} \cdot d \leq 8d^2 \cdot \log h \leq \frac{h}{4} \cdot \log \frac{6}{5}.$$

Re-arrange the inequality and notice that $h - d \leq \gamma$, we have

$$2d^2 + 3d + 1 \leq \frac{(h - d) \log \frac{6}{5}}{4 \log(h - d)} \leq \frac{\gamma \log \frac{6}{5}}{4 \log \gamma}, \quad (10)$$

therefore,

$$\gamma^{2d^2+3d+1} \leq \left(\frac{6}{5}\right)^{\frac{\gamma}{4}}. \quad (11)$$

Next, note that for any integers a, b such that $a < b$, we have $\frac{\binom{\gamma}{a}}{\binom{\gamma}{b}} = \frac{(a+1)(a+2)\cdots b}{(\gamma-b+1)(\gamma-b+2)\cdots(\gamma-a)} \leq \left(\frac{b}{\gamma-a}\right)^{b-a}$. Because $z^* \in [\lceil \frac{\gamma}{2} \rceil, \lceil \frac{\gamma}{2} \rceil + 2d^2]$, we have $\binom{\gamma}{z^*} \geq \binom{\gamma}{\lceil \frac{\gamma}{2} \rceil + 2d^2}$. Thus,

$$\begin{aligned} \frac{\binom{\gamma}{t}}{\binom{\gamma}{z^*}} &\leq \frac{\binom{\gamma}{t}}{\binom{\gamma}{\lceil \frac{\gamma}{2} \rceil + 2d^2}} \leq \left(\frac{\lceil \frac{\gamma}{2} \rceil + 2d^2}{\gamma - t}\right)^{\lceil \frac{\gamma}{2} \rceil + 2d^2 - t} \\ &\leq \left(\frac{\frac{\gamma}{2} + 2d^2 + 1}{\frac{3}{4}\gamma}\right)^{\lceil \frac{\gamma}{2} \rceil + 2d^2 - t}. \end{aligned} \quad (12)$$

Because $\gamma \geq h - d$, for any $h \geq 2$ and $d \leq \frac{1}{8}\sqrt{\frac{h}{2\log h}}$, it must be $\frac{\log \frac{6}{5}}{\log \gamma} < \frac{1}{2}$. By Equation (10), we have

$$2d^2 + 1 \leq 2d^2 + 3d + 1 \leq \frac{\gamma \log \frac{6}{5}}{4 \log \gamma} \leq \frac{\gamma}{8}.$$

By $2d^2 + 1 \leq \frac{\gamma}{8}$ and Equation (12), we have

$$\frac{\binom{\gamma}{t}}{\binom{\gamma}{z^*}} \leq \left(\frac{\frac{\gamma}{2} + 2d^2 + 1}{\frac{3}{4}\gamma}\right)^{\lceil \frac{\gamma}{2} \rceil + 2d^2 - t} \leq \left(\frac{5}{6}\right)^{\lceil \frac{\gamma}{2} \rceil + 2d^2 - t} \leq \left(\frac{5}{6}\right)^{\frac{\gamma}{4}}. \quad (13)$$

Combining Equations (9), (11), and (13), we have that $\|\delta\|_\infty \leq 1$. \square

5 Characterization for Finite Block Size

5.1 Characterization for Strict IC

In this section, we give a characterization for strict incentive compatibility for finite block size. In an (h, ρ, c, d) -environment, we can indeed circumvent the 0-miner revenue impossibility result in [SCW23]. However, it turns out that for $c = 1$ and $c \geq 2$, the mechanisms are different. Specifically, for $c \geq 2$, each user's utility has to be 0. Therefore, we separately give the mechanisms for $c = 1$ and $c \geq 2$.

5.1.1 Feasibility for $c = 1$

For $c = 1$, the mechanism is simply the LP-based mechanism in Section 4.2 with a random selection process. Still, we assume that honest users' values are sampled i.i.d. from some distribution \mathcal{D} , and the median m of the distribution satisfies that $\Pr[x \geq m] = \frac{1}{2}$ (see Remark 2.1). For convenience, we repeat the MPC-assisted, LP-based mechanism with random selection, which has been introduced in Section 2.3.1.

MPC-assisted, LP-based mechanism with random selection

Parameters: the block size k , the environment parameter $(h, *, 1, d)$, the distribution median m .

Input: a bid vector $\mathbf{b} = (b_1, \dots, b_n)$.

Mechanism:

- *Confirmation Rule.* Let $\tilde{\mathbf{b}} = (\tilde{b}_1, \dots, \tilde{b}_s)$ denote the bids that are at least m . If $s \leq k$, confirm all bids in $\tilde{\mathbf{b}}$. Otherwise, randomly select k bids from $\tilde{\mathbf{b}}$ to confirm.
- *Payment rule.* Each confirmed bid pays m .
- *Miner revenue rule.* Let $\mathbf{y} := (y_0, y_1, \dots, y_n)$ be any feasible solution to the following linear program:

$$\forall i \in [n] : 0 \leq y_i \leq \min(i, k) \cdot m \quad (14)$$

$$\forall 0 \leq j \leq d : \sum_{i=0}^{n-d} q_i \cdot y_{i+j} = \frac{m \cdot \min(h, k)}{4} \quad (15)$$

where $q_i = \frac{1}{2^{n-d}} \binom{n-d}{i}$ is the probability of observing i heads if we flip $n - d$ independent fair coins. The total miner revenue is y_t where t is the number of confirmed bids in the block.

Theorem 5.1. *Suppose the block size is k . Fix any⁸ $h \geq 2$, and any $d \leq \frac{1}{8} \sqrt{\frac{h}{2 \log h}}$. The MPC-assisted, LP-based mechanism with random selection is ex post UIC, Bayesian MIC, and Bayesian SCP in an $(h, *, 1, d)$ -environment. Moreover, the expected miner revenue is $\Theta(\min\{h, k\})$.*

Proof. First, Equation (14) guarantees that total miner revenue is at most the total payment of the confirmed users, so the mechanism satisfies budget feasibility.

⁸For the special case $h = 1$, we can just use the parity-based mechanism of Section 2.2 with the random selection.

Next, we show that when the linear program Equations (14) and (15) has a solution, the mechanism satisfies all three incentive-compatible properties.

- **UIC:** By the same reasoning as in Theorem A.1, overbidding or underbidding does not increase the user's utility. Injecting bids cannot increase the user's utility either: it may only decrease the probability that the user gets confirmed. Moreover, dropping out can only give the user zero utility. Therefore, a user cannot increase its utility by deviating.
- **SCP:** By the same reasoning as in the proof of Lemma 4.6, the linear program Equation (15) guarantees that no matter how the coalition chooses the d bids it controls, the expected total miner revenue remains unchanged. Meanwhile, the coalition cannot increase the colluding user's utility by UIC. Therefore, this mechanism is SCP.
- **MIC:** Follows by the same reasoning as SCP.

It remains to show that the linear program indeed has a feasible solution. We will give a constructive solution. Let $\tilde{\mathbf{y}} = (\tilde{y}_1, \dots, \tilde{y}_n)$ denote the constructive solution given in the proof of Lemma 4.7 that satisfies

$$\forall 0 \leq j \leq d : \sum_{i=0}^{n-d} q_i \cdot \tilde{y}_{i+j} = \frac{m \cdot h}{4}.$$

In the proof of Lemma 4.7, $\tilde{\mathbf{y}}$ satisfies that $0 \leq \tilde{y}_i \leq \min(i, h) \cdot m$ for any $0 \leq i \leq n$. There are two possible cases.

- $h \leq k$. We have $0 \leq \tilde{y}_i \leq \min(i, h) \cdot m \leq \min(i, k) \cdot m$. Thus, $\tilde{\mathbf{y}}$ is a feasible solution to the linear program in this case.
- $h > k$. Let $\mathbf{y} = (y_0, \dots, y_n) = \frac{k}{h} \cdot \tilde{\mathbf{y}}$. Then y_i satisfies that $0 \leq y_i \leq \frac{k}{h} \min(i, h) \cdot m \leq \min(i, k) \cdot m$. Moreover, for any $0 \leq j \leq d$,

$$\sum_{i=0}^{n-d} q_i \cdot y_{i+j} = \frac{k}{h} \cdot \sum_{i=0}^{n-d} q_i \cdot \tilde{y}_{i+j} = \frac{m \cdot k}{4}.$$

Thus, $\mathbf{y} = \frac{k}{h} \cdot \tilde{\mathbf{y}}$ is a feasible solution to the linear program if $h > k$.

□

5.1.2 Zero Social Welfare for Users When $c \geq 2$

Unfortunately, the above MPC-assisted, LP-based mechanism with random selection only works for $c = 1$. When $c \geq 2$, although deviating cannot increase the expected total miner revenue, the coalition can increase a colluding user's utility. Imagine that the coalition consists of some colluding miners and two users i and j , where user i has true value m and user j has a large true value. Then user i may choose to drop out to increase the probability of user j getting confirmed. This strictly increases the expected joint utility of the coalition.

Therefore, to construct a Bayesian SCP mechanism in an (h, ρ, c, d) -environment for $d \geq c \geq 2$, we need to make sure that deviating cannot increase a colluding user's utility. Indeed, for some (contrived) distributions, we can construct a mechanism that generates optimal miner revenue and achieves UIC, MIC, and SCP in an (h, ρ, c, d) -environment for $d \geq c \geq 2$. However, the total social welfare for all users is 0. For example, imagine that honest users' true values are drawn i.i.d. from

Bernoulli($\frac{1}{2}$). Now, if we run the MPC-assisted, LP-based mechanism with random selection (see Section 5.1.1) and set $m = 1$, the resulting mechanism achieves ex post UIC, Bayesian MIC, and Bayesian SCP in $(h, *, c, d)$ -environments, even when $c \geq 2$ (as long as the condition $d \leq \frac{1}{8} \sqrt{\frac{h}{2 \log h}}$ is satisfied). This is because setting $m = 1$ makes sure that every user's utility is always 0. Thus, no matter how the coalition deviates, it cannot increase the strategic users' joint utility. Moreover, as long as the linear program Equations (14) and (15) has a feasible solution, the coalition cannot increase the expected total miner revenue either. The mechanism achieves $\Theta(m)$ expected miner revenue but unfortunately, the total user social welfare is always 0. It turns out that this zero user social welfare limitation is intrinsic, as stated below.

Theorem 5.2 (Restatement of Theorem 1.6). *Suppose that the block size is finite, and fix any $h \geq 1$, any $d \geq c \geq 2$, and any $\rho \in (0, 1)$. Then, any MPC-assisted TFM that simultaneously satisfies Bayesian UIC, MIC and SCP in an (h, ρ, c, d) -model must suffer from 0 social welfare for the users when there actually are more than h honest bids. Equivalently, for any $\ell > h$,*

$$\mathbf{E}_{\mathbf{b} \sim \mathcal{D}^\ell} [\text{USW}(\mathbf{b})] = 0. \quad (16)$$

In the above, $\text{USW}(\mathbf{b})$ denotes the expected total user social welfare under the bid vector \mathbf{b} where the expectation is taken over the randomness of the mechanism.

The proof is similar to the proof of Theorem 5.2 of [SCW23]. We will use the following lemma of [SCW23] to prove this theorem. Although the original lemma considers a universal MPC-assisted mechanism, the proof also holds for MPC-assisted TFM in an (h, ρ, c, d) -environment for $d \geq c \geq 2$. Henceforth, we use $\text{util}^i(\mathbf{b})$ to denote the utility of identity i when the input bid vector is \mathbf{b} . In the proof, we use $v_{\text{id}}(b_{\text{id}})$ to denote a bid v (b) coming from identity id .

Lemma 5.3. *Fix any $h \geq 1$, any $d \geq c \geq 2$, any $\rho \in (0, 1)$. Given any (possibly random) MPC-assisted mechanism that is Bayesian UIC, MIC and SCP in an (h, ρ, c, d) -environment, for any identity i and identity j , for any bid b_j and b'_j , it must be that for any $\ell \geq h$,*

$$\mathbf{E}_{(v_i, \mathbf{b}_{-i,j}) \sim \mathcal{D}^{\ell+1}} [\text{util}^i(v_i, b_j, \mathbf{b}_{-i,j})] = \mathbf{E}_{(v_i, \mathbf{b}_{-i,j}) \sim \mathcal{D}^{\ell+1}} [\text{util}^i(v_i, b'_j, \mathbf{b}_{-i,j})], \quad (17)$$

where $\mathbf{b}_{-i,j}$ represents all except identity i and j 's bids. Moreover, it must be that

$$\mathbf{E}_{(v_i, \mathbf{b}_{-i,j}) \sim \mathcal{D}^{\ell+1}} [\text{util}^i(v_i, b_j, \mathbf{b}_{-i,j})] = \mathbf{E}_{(v_i, \mathbf{b}_{-i}) \sim \mathcal{D}^{\ell+1}} [\text{util}^i(v_i, \mathbf{b}_{-i})]. \quad (18)$$

Proof. The proof to this lemma is the same as in Lemma 5.2 and 5.3 of [SCW23], except that now we need to guarantee that at least h bids are sampled randomly from \mathcal{D} . \square

Corollary 5.4. *Fix any $h \geq 1$, any $d \geq c \geq 2$, any $\rho \in (0, 1)$. Given any (possibly random) MPC-assisted mechanism that is Bayesian UIC, MIC and SCP in an (h, ρ, c, d) -environment, for any two sets H and H' consisting of at least h identities, let \mathbf{b}_H ($\mathbf{b}_{H'}$) denote the bids from identities in H (H'). For any $i \notin H \cup H'$, it must be that*

$$\mathbf{E}_{(v_i, \mathbf{b}_H) \sim \mathcal{D}^{|H|+1}} [\text{util}^i(v_i, \mathbf{b}_H)] = \mathbf{E}_{(v_i, \mathbf{b}_{H'}) \sim \mathcal{D}^{|H'|+1}} [\text{util}^i(v_i, \mathbf{b}_{H'})],$$

where v_i denotes that identity i bids v .

Proof. Let $S = H' \setminus H$. Without loss of generality, we assume that S consists of identities $1, \dots, |S|$. By the definition of the total expectation, we have

$$\begin{aligned}
& \mathbf{E}_{(v_i, \mathbf{b}_S, \mathbf{b}_H) \sim \mathcal{D}^{|S|+|H|+1}} [\text{util}^i(v_i, \mathbf{b}_S, \mathbf{b}_H)] \\
&= \int_0^\infty \mathbf{E}_{(v_i, \mathbf{b}_{S \setminus \{1\}}, \mathbf{b}_H) \sim \mathcal{D}^{|S|+|H|}} [\text{util}^i(v_i, z_1, \mathbf{b}_{S \setminus \{1\}}, \mathbf{b}_H)] f(z_1) dz_1 \\
&= \mathbf{E}_{(v_i, \mathbf{b}_{S \setminus \{1\}}, \mathbf{b}_H) \sim \mathcal{D}^{|S|+|H|}} [\text{util}^i(v_i, b_1, \mathbf{b}_{S \setminus \{1\}}, \mathbf{b}_H)] \int_0^\infty f(z_1) dz_1 && \text{By Equation (17)} \\
&= \mathbf{E}_{(v_i, \mathbf{b}_{S \setminus \{1\}}, \mathbf{b}_H) \sim \mathcal{D}^{|S|+|H|}} [\text{util}^i(v_i, b_1, \mathbf{b}_{S \setminus \{1\}}, \mathbf{b}_H)] \\
&= \mathbf{E}_{(v_i, \mathbf{b}_{S \setminus \{1\}}, \mathbf{b}_H) \sim \mathcal{D}^{|S|+|H|}} [\text{util}^i(v_i, \mathbf{b}_{S \setminus \{1\}}, \mathbf{b}_H)] && \text{By Equation (18)} \\
&= \dots = \mathbf{E}_{(v_i, \mathbf{b}_H) \sim \mathcal{D}^{|H|+1}} [\text{util}^i(v_i, \mathbf{b}_H)].
\end{aligned}$$

By the same reasoning, consider $S' = H \setminus H'$. Then it must be that

$$\mathbf{E}_{(v_i, \mathbf{b}_{S'}, \mathbf{b}_{H'}) \sim \mathcal{D}^{|S'|+|H'|+1}} [\text{util}^i(v_i, \mathbf{b}_{S'}, \mathbf{b}_{H'})] = \mathbf{E}_{(v_i, \mathbf{b}_{H'}) \sim \mathcal{D}^{|H'|+1}} [\text{util}^i(v_i, \mathbf{b}_{H'})].$$

Note that $S' \cup H' = S \cup H = H' \cup H$. Hence,

$$\mathbf{E}_{(v_i, \mathbf{b}_{S'}, \mathbf{b}_{H'}) \sim \mathcal{D}^{|S'|+|H'|+1}} [\text{util}^i(v_i, \mathbf{b}_{S'}, \mathbf{b}_{H'})] = \mathbf{E}_{(v_i, \mathbf{b}_S, \mathbf{b}_H) \sim \mathcal{D}^{|S|+|H|+1}} [\text{util}^i(v_i, \mathbf{b}_S, \mathbf{b}_H)].$$

Combining the equalities, we have

$$\begin{aligned}
& \mathbf{E}_{(v_i, \mathbf{b}_H) \sim \mathcal{D}^{|H|+1}} [\text{util}^i(v_i, \mathbf{b}_H)] \\
&= \mathbf{E}_{(v_i, \mathbf{b}_S, \mathbf{b}_H) \sim \mathcal{D}^{|S|+|H|+1}} [\text{util}^i(v_i, \mathbf{b}_S, \mathbf{b}_H)] \\
&= \mathbf{E}_{(v_i, \mathbf{b}_{S'}, \mathbf{b}_{H'}) \sim \mathcal{D}^{|S'|+|H'|+1}} [\text{util}^i(v_i, \mathbf{b}_{S'}, \mathbf{b}_{H'})] \\
&= \mathbf{E}_{(v_i, \mathbf{b}_{H'}) \sim \mathcal{D}^{|H'|+1}} [\text{util}^i(v_i, \mathbf{b}_{H'})]
\end{aligned}$$

□

This corollary implies that when identity i 's bid is sampled from \mathcal{D} in a world with h or more random bids, its expected utility only depends on its identity i . Henceforth we will use the following notation to denote this utility (where the notation v_i means identity i is bidding the value v):

$$U_i := \mathbf{E}_{(v_i, \mathbf{b}_H) \sim \mathcal{D}^{|H|+1}} [\text{util}^i(v_i, \mathbf{b}_H)]. \tag{19}$$

Lemma 5.5. *Fix any $h \geq 1$, any $d \geq c \geq 2$, any $\rho \in (0, 1)$. Given any (possibly random) MPC-assisted mechanism that is Bayesian UIC, MIC and SCP in an (h, ρ, c, d) -environment, for any user i, j , it must be that*

$$U_i = U_j. \tag{20}$$

Proof. Fix any set H of at least $h + 1$ number of users. By our symmetric assumption, it must be that

$$\mathbf{E}_{\mathbf{b}_H \sim \mathcal{D}^{|H|}} [\text{USW}(v_i, \mathbf{b}_H)] = \mathbf{E}_{\mathbf{b}_H \sim \mathcal{D}^{|H|}} [\text{USW}(v_j, \mathbf{b}_H)], \quad (21)$$

where v_i (v_j) denotes that identity i (j) bids v , and $\text{USW}(\mathbf{b})$ denotes the expected social welfare for all users when the input bid vector is \mathbf{b} . For any identity $l \in H$, for any v_i from identity i , let $H' = H \setminus \{l\}$. It must be

$$\begin{aligned} \mathbf{E}_{(b_l, \mathbf{b}_{H'}) \sim \mathcal{D}^{|H|}} [\text{util}^l(b_l, v_i, \mathbf{b}_{H'})] &= \mathbf{E}_{(b_l, \mathbf{b}_{H'}) \sim \mathcal{D}^{|H|}} [\text{util}^l(b_l, \mathbf{b}_{H'})] && \text{By Equation (18)} \\ &= U_l && \text{By Equation (19)} \end{aligned}$$

By the same reasoning, $\mathbf{E}_{(b_l, \mathbf{b}_{H'}) \sim \mathcal{D}^{|H|}} [\text{util}^l(b_l, v_j, \mathbf{b}_{H'})] = U_l$. Thus, for any value v , the sum of the expected utility of every user in H is

$$\sum_{l \in H} \mathbf{E}_{\mathbf{b}_H \sim \mathcal{D}^{|H|}} [\text{util}^l(v_i, \mathbf{b}_H)] = \sum_{l \in H} U_l = \sum_{l \in H} \mathbf{E}_{\mathbf{b}_H \sim \mathcal{D}^{|H|}} [\text{util}^l(v_j, \mathbf{b}_H)].$$

Combining this with Equation (21), it must be that for any v_i and v_j (which denote that identity i and j bid value v , respectively),

$$\mathbf{E}_{\mathbf{b}_H \sim \mathcal{D}^{|H|}} [\text{util}^i(v_i, \mathbf{b}_H)] = \mathbf{E}_{\mathbf{b}_H \sim \mathcal{D}^{|H|}} [\text{util}^j(v_j, \mathbf{b}_H)].$$

The lemma follows by taking expectations over v on both sides. \square

Lemma 5.6. *Fix any $h \geq 1$, any $d \geq c \geq 2$, any $\rho \in (0, 1)$, and suppose that the distribution \mathcal{D} has bounded support. Given any (possibly random) MPC-assisted mechanism that is Bayesian UIC, MIC and SCP in an (h, ρ, c, d) -environment, for any identity i , it must be that*

$$U_i = 0.$$

Proof. Consider a crowded world with many users and all of their bids are sampled independently at random from \mathcal{D} . Let K be the total number of users. By Corollary 5.4 and Lemma 5.5, every user's expected utility is the same where the expectation is taken over the random coins for sampling all bids as well as random coins of the mechanism. On the other hand, since there are K total bids, there must exist a user whose confirmation probability is at most k/K , and thus its expected utility is at most $\max(\mathcal{D}) \cdot k/K$ where k is the block size. The lemma follows by taking K to be arbitrarily large. \square

Proof of Theorem 5.2 Fix any set H of size at least $h + 1$. Then

$$\mathbf{E}_{\mathbf{b} \in \mathcal{D}^{|H|}} [\text{USW}(\mathbf{b})] = \sum_{i \in H} \mathbf{E}_{\mathbf{b} \in \mathcal{D}^{|H|}} \text{util}^i(\mathbf{b}).$$

By Lemma 5.6, for each identity i in H ,

$$\mathbf{E}_{\mathbf{b} \in \mathcal{D}^{|H|}} \text{util}^i(\mathbf{b}) = U_i = 0.$$

Therefore, the user social welfare is 0.

5.2 Feasibility for Approximate IC: Diluted Threshold-Based Mechanism

Although there is no interesting mechanism for strict incentive compatibility when $c \geq 2$, there are meaningful mechanisms if we allow approximate incentive compatibility. Still, we assume that honest users' values are sampled i.i.d. from some bounded distribution \mathcal{D} , and m is the median of \mathcal{D} such that $\Pr[x \geq m] = \frac{1}{2}$ (see Remark 2.1). In addition, we assume that there is an upper bound T on users' true values: $\Pr[x \leq T] = 1$. Without loss of generality, we assume $T \geq \epsilon$.

MPC-assisted, diluted threshold-based Mechanism

Parameters: the block size k , the environment parameter $(h, *, c, *)$, the approximation parameter ϵ , the distribution median m , and the upper bound T of the distribution.

Input: a bid vector $\mathbf{b} = (b_1, \dots, b_n)$.

Mechanism:

- *Confirmation rule.* Let $R := \max\left(2c\sqrt{\frac{kT}{\epsilon}}, k\right)$. Given a bid vector \mathbf{b} , let $\tilde{\mathbf{b}} = (\tilde{b}_1, \dots, \tilde{b}_s)$ denote the bids that are at least m . If $s \leq R$, randomly select $\frac{k}{R} \cdot s$ bids from $\tilde{\mathbf{b}}$ to confirm; otherwise, randomly select k bids from $\tilde{\mathbf{b}}$ to confirm.
- *Payment rule.* Every confirmed bid pays m .
- *Miner revenue rule.* If $s \geq \frac{h}{4}$, the total miner revenue is $\bar{\mu} := m \cdot \min\left(\frac{h}{4} \cdot \frac{k}{R}, k\right)$. Otherwise, the total miner revenue is 0.

Theorem 5.7. *Suppose the block size is k . For any $h \geq 1$, $c \geq 1$, and $\epsilon \geq m \cdot \frac{h}{2} \cdot e^{-\frac{h}{16}}$, the diluted threshold posted price auction satisfies strict ex post UIC, Bayesian ϵ -MIC, and Bayesian ϵ -SCP in an $(h, *, c, *)$ -environment. Moreover, the expected total miner revenue is $m \cdot \min\left(\frac{h\sqrt{k\epsilon}}{8c\sqrt{T}}, \frac{h}{4}, k\right)$, where T is the upper bound of users' true values.*

Proof. We first show that the budget feasibility is satisfied. Since the mechanism confirms $\min\left(s \cdot \frac{k}{R}, k\right)$ number of bids that are at least m , the total payment is $m \cdot \min\left(s \cdot \frac{k}{R}, k\right)$. When $s \geq \frac{h}{4}$, the total miner revenue is at most $m \cdot \frac{h}{4} \cdot \frac{k}{R} \leq m \cdot s \cdot \frac{k}{R}$, which is no more than the total payment of the users. Next, we prove UIC, MIC, and SCP separately.

UIC. Since the mechanism is posted price auction from each user's perspective, each user's best response is to follow the protocol honestly, as in the proof of Theorem 5.1.

MIC. By the same reasoning as in Theorem 4.1, by injecting fake bids, the miner can only increase its expected miner revenue if the number of bids that are at least m from honest users is less than $\frac{h}{4}$. This happens with a probability at most $e^{-\frac{h}{16}}$. Thus, the expected total miner revenue increases by no more than

$$\bar{\mu} \cdot e^{-\frac{h}{16}} \leq m \cdot e^{-\frac{h}{16}} \cdot \frac{h}{4} \leq \frac{\epsilon}{2}.$$

SCP. By the same reasoning as in MIC, the expected increase of the miner revenue is at most $\epsilon/2$ by any deviation. Thus, to show that the mechanism is Bayesian ϵ -SCP, it suffices to show that the coalition cannot increase the joint utility of the "users" in the coalition by more than $\frac{\epsilon}{2}$.

Because injecting bids smaller than m does not change the confirmation probability and the payment of each confirmed bid is fixed, injecting bids smaller than m does not increase the users' utilities. On the other hand, injecting bids at least m will only decrease the probability of each colluding user getting confirmed, which does not increase the users' utilities.

Now, it suffices to show that overbidding and underbidding do not increase the coalition's joint utility since dropping out is equivalent to underbidding to some value less than m . Let s be the number of bids $\geq m$ when every user bids truthfully. Each bid is confirmed with probability $\frac{k}{R}$ if $s \leq R$, and $\frac{k}{s}$ if $s > R$. Let s' be the number of bids $\geq m$ when the colluding users bid strategically. The colluding users can be partitioned into four groups:

- S_1 : Those whose true values are less than m but overbid to values larger than or equal to m ;
- S_2 : Those whose true values are less than m and bid values less than m ;
- S_3 : Those whose true values are at least m but underbids to values less than m ;
- S_4 : Those whose true values are at least m and still bid values at least m .

When the coalition bids strategically, only the utilities of the users in S_4 increase compared to the honest case. Consider a colluding user in S_4 with the true value $v \geq m$. Its utility increases by at most

$$(v - m) \cdot \frac{k}{\max\{s', R\}} - (v - m) \cdot \frac{k}{\max\{s, R\}}. \quad (22)$$

Note that Equation (22) is positive only when $s' < s$ and $s > R$. In this case, Equation (22) can be upper bounded by

$$\begin{aligned} (v - m) \left[\frac{k}{\max\{s', R\}} - \frac{k}{s} \right] &\leq (T - m) \left[\frac{k}{s'} - \frac{k}{s} \right] \\ &\leq (T - m) \cdot \frac{ck}{s(s - c)} && \text{By } s' \geq s - c. \\ &\leq T \cdot \frac{ck}{R(R - c)}. \end{aligned}$$

Since by the choice of R , $R(R - c) \geq \frac{1}{2}R^2$, we have

$$\text{Equation (22)} \leq T \cdot \frac{2ck}{R^2} \leq \frac{\epsilon}{2c}.$$

Therefore, by bidding untruthfully, each user's utility can increase by at most $\frac{\epsilon}{2c}$. Therefore, the joint utility of the users in the coalition by more than $\frac{\epsilon}{2}$. \square

6 Bounds on Miner Revenue

In this section, we prove bounds on the miner revenue under different settings. Henceforth, let $\mu(\mathbf{b})$ denote the expected total miner revenue when the input bid vector is \mathbf{b} , where the expectation is taken over the mechanism's randomness.

6.1 Known- h Model

In this section, we prove limits on miner revenue in the known- h model (i.e., Theorem 1.1).

Theorem 6.1 (Limit on miner revenue for approximate incentive compatibility, Theorem 1.1 restated). *Fix any $h \geq 1$, $d \geq c \geq 1$, and $\rho \in (0, 1)$. Given any MPC-assisted mechanism that is Bayesian ϵ_u -UIC, Bayesian ϵ_m -MIC and Bayesian ϵ_s -SCP in an (h, ρ, c, d) -environment, for all $n \geq h$, it must be that*

$$\mathbf{E}_{\mathbf{b} \sim \mathcal{D}^n} [\mu(\mathbf{b})] \leq h \cdot \mathbf{E}(\mathcal{D}) + \frac{2(n-h)}{\rho} (\epsilon + C_{\mathcal{D}} \sqrt{\epsilon}), \quad (23)$$

where $\epsilon = \epsilon_u + \epsilon_m + \epsilon_s$, $\mathbf{E}(\mathcal{D}) = \mathbf{E}_{X \sim \mathcal{D}}[X]$ and $C_{\mathcal{D}} = \mathbf{E}_{X \sim \mathcal{D}}[\sqrt{X}]$ are the terms that depend on the “scale” of the distribution \mathcal{D} .

As a special case, for strict incentive compatibility where $\epsilon = 0$, we have that

$$\mathbf{E}_{\mathbf{b} \sim \mathcal{D}^n} [\mu(\mathbf{b})] \leq h \cdot \mathbf{E}(\mathcal{D}) \quad (24)$$

To prove Theorem 6.1, we need the following lemma.

Lemma 6.2 (Lemma 3.3 of [SCW23]). *Fix any $h \geq 1$, $d \geq c \geq 1$ and $\rho \in (0, 1)$. Given any (possibly randomized) MPC-assisted TFM that is Bayesian ϵ_u -UIC and Bayesian ϵ_s -SCP in an (h, ρ, c, d) -environment, for any user i and any value v , for any $\ell \geq h$ it must be that*

$$\mathbf{E}_{\mathbf{b} \sim \mathcal{D}^\ell} [\mu(\mathbf{b}, v)] - \mathbf{E}_{\mathbf{b} \sim \mathcal{D}^\ell} [\mu(\mathbf{b}, 0)] \leq \begin{cases} \frac{2}{\rho}(\epsilon_s + \epsilon_u), & \text{if } v \leq \epsilon_s + \epsilon_u; \\ \frac{2}{\rho}(\sqrt{v(\epsilon_s + \epsilon_u)}), & \text{if } v > \epsilon_s + \epsilon_u. \end{cases} \quad (25)$$

Proof. [SCW23] stated for any $\rho \in (0, 1)$ and any $c \geq 1$, given any mechanism that is Bayesian ϵ_u -UIC and Bayesian ϵ_s -SCP in an $(*, \rho, c, *)$ -environment, for any user i and any value v , for any $\ell \geq 0$, Equation (25) must hold.

In their proof, they showed that if for some ℓ and some v , Equation (25) is violated, and moreover assuming that the mechanism satisfies Bayesian ϵ_s -SCP, then there exists some value v' such that if a user of value v' colludes with a subset of the miners, the coalition can play strategically and jointly benefit when the rest of the world contains ℓ honest users. We can apply the exactly same argument, but because it is promised that there are at least h honest users, the argument only holds for $\ell \geq h$. Observe also the coalition’s strategy requires only $d = 1$. □

Now, we are ready to prove Theorem 6.1.

Proof of Theorem 6.1. The proof mainly follows the proof of Theorem 3.4 in [SCW23]. Since the TFM is Bayesian ϵ_m -MIC in an (h, ρ, c, d) -environment, it must be that for any $\ell \geq h$,

$$\mathbf{E}_{\mathbf{b} \sim \mathcal{D}^\ell} [\rho \mu(\mathbf{b}, 0)] \leq \mathbf{E}_{\mathbf{b} \sim \mathcal{D}^\ell} [\rho \mu(\mathbf{b})] + \epsilon_m. \quad (26)$$

Otherwise, when there are ℓ honest bids, a strategic ρ -sized miner coalition can inject a bid 0 and increase its miner revenue by strictly more than ϵ_m , while it does not need to pay anything for injecting this 0-bid. This violates Bayesian ϵ_m -MIC.

Let $f(\cdot)$ be the p.d.f. of distribution \mathcal{D} . By the law of total expectation, for all $\ell \geq 1$, we have

$$\mathbf{E}_{\mathbf{b} \sim \mathcal{D}^\ell} [\mu(\mathbf{b})] = \int_0^\infty \mathbf{E}_{\mathbf{b}' \sim \mathcal{D}^{\ell-1}} [\mu(\mathbf{b}', r)] f(r) dr.$$

Let $\epsilon' = \epsilon_s + \epsilon_u$. Since the mechanism is Bayesian ϵ_u -UIC and Bayesian ϵ_s -SCP in an (h, ρ, c, d) -environment, by Lemma 6.2, for all $\ell \geq h + 1$, it must be that

$$\begin{aligned} \int_0^{\epsilon'} \mathbf{E}_{\mathbf{b}' \sim \mathcal{D}^{\ell-1}} [\mu(\mathbf{b}', r)] f(r) dr &\leq \int_0^{\epsilon'} \left[\mathbf{E}_{\mathbf{b}' \sim \mathcal{D}^{\ell-1}} [\mu(\mathbf{b}', 0)] + \frac{2\epsilon'}{\rho} \right] f(r) dr; \\ \int_{\epsilon'}^{\infty} \mathbf{E}_{\mathbf{b}' \sim \mathcal{D}^{\ell-1}} [\mu(\mathbf{b}', r)] f(r) dr &\leq \int_{\epsilon'}^{\infty} \left[\mathbf{E}_{\mathbf{b}' \sim \mathcal{D}^{\ell-1}} [\mu(\mathbf{b}', 0)] + \frac{2\sqrt{r\epsilon'}}{\rho} \right] f(r) dr. \end{aligned}$$

Summing up the two inequalities above, we can bound the expected miner revenue with

$$\begin{aligned} &\mathbf{E}_{\mathbf{b} \sim \mathcal{D}^\ell} [\mu(\mathbf{b})] \\ &= \int_0^{\epsilon'} \mathbf{E}_{\mathbf{b}' \sim \mathcal{D}^{\ell-1}} [\mu(\mathbf{b}', r)] f(r) dr + \int_{\epsilon'}^{\infty} \mathbf{E}_{\mathbf{b}' \sim \mathcal{D}^{\ell-1}} [\mu(\mathbf{b}', r)] f(r) dr \\ &\leq \int_0^{\epsilon'} \left[\mathbf{E}_{\mathbf{b}' \sim \mathcal{D}^{\ell-1}} [\mu(\mathbf{b}', 0)] + \frac{2\epsilon'}{\rho} \right] f(r) dr \\ &\quad + \int_{\epsilon'}^{\infty} \left[\mathbf{E}_{\mathbf{b}' \sim \mathcal{D}^{\ell-1}} [\mu(\mathbf{b}', 0)] + \frac{2\sqrt{r\epsilon'}}{\rho} \right] f(r) dr \\ &\leq \mathbf{E}_{\mathbf{b}' \sim \mathcal{D}^{\ell-1}} [\mu(\mathbf{b}', 0)] + \frac{2\epsilon'}{\rho} \int_0^{\epsilon'} f(r) dr + \frac{2\sqrt{\epsilon'}}{\rho} \int_{\epsilon'}^{\infty} \sqrt{r} f(r) dr \end{aligned}$$

By (26), we have that $\mathbf{E}_{\mathbf{b}' \sim \mathcal{D}^{\ell-1}} [\mu(\mathbf{b}', 0)] \leq \mathbf{E}_{\mathbf{b}' \sim \mathcal{D}^{\ell-1}} [\mu(\mathbf{b}')] + \frac{\epsilon_s}{\rho}$. Therefore, for all $\ell \geq h + 1$,

$$\begin{aligned} &\mathbf{E}_{\mathbf{b} \sim \mathcal{D}^\ell} [\mu(\mathbf{b})] \\ &\leq \mathbf{E}_{\mathbf{b}' \sim \mathcal{D}^{\ell-1}} [\mu(\mathbf{b}', 0)] + \frac{2\epsilon'}{\rho} \int_0^{\epsilon'} f(r) dr + \frac{2\sqrt{\epsilon'}}{\rho} \int_{\epsilon'}^{\infty} \sqrt{r} f(r) dr \\ &\leq \mathbf{E}_{\mathbf{b}' \sim \mathcal{D}^{\ell-1}} [\mu(\mathbf{b}')] + \frac{\epsilon_s}{\rho} + \frac{2\epsilon'}{\rho} + \frac{2\sqrt{\epsilon'}}{\rho} \mathbf{E}_{X \sim \mathcal{D}} [\sqrt{X}] \\ &\leq \mathbf{E}_{\mathbf{b}' \sim \mathcal{D}^{\ell-1}} [\mu(\mathbf{b}')] + \frac{2\epsilon}{\rho} + \frac{2C_{\mathcal{D}}\sqrt{\epsilon}}{\rho}, \end{aligned} \tag{27}$$

where the last step comes from the fact that $\epsilon = \epsilon_u + \epsilon_m + \epsilon_s$.

Finally, for any bid vector \mathbf{b} such that at least h bids submitted by honest users, we can represent \mathbf{b} as $\mathbf{b} = \mathbf{b}_H + \mathbf{b}_{-H}$, where \mathbf{b}_H are submitted all from honest users and $|\mathbf{b}_H| = h$. Notice that \mathbf{b}_{-H} might or might not contain the bids submitted by honest users. With this notation, for all $\ell \geq h$, we can write

$$\mathbf{E}_{\mathbf{b} \sim \mathcal{D}^\ell} [\mu(\mathbf{b})] = \mathbf{E}_{\mathbf{b}_H \sim \mathcal{D}^h, \mathbf{b}_{-H} \sim \mathcal{D}^{\ell-h}} [\mu(\mathbf{b}_H + \mathbf{b}_{-H})].$$

By applying Equation (27) $(\ell - h)$ times, we have

$$\mathbf{E}_{\mathbf{b} \sim \mathcal{D}^\ell} [\mu(\mathbf{b})] \leq \mathbf{E}_{\mathbf{b}_H \sim \mathcal{D}^h} [\mu(\mathbf{b}_H)] + \frac{2(\ell - h)}{\rho} \left(\epsilon + \sqrt{\epsilon} \cdot \mathbf{E}_{X \sim \mathcal{D}} [\sqrt{X}] \right).$$

Because honest users always submit their true values and the miner revenue is bounded by the sum of the payments, if the bids are submitted by h honest users, it must be

$$\mathbf{E}_{\mathbf{b} \sim \mathcal{D}^h} [\mu(\mathbf{b})] \leq h \cdot \mathbf{E}_{X \sim \mathcal{D}} [X].$$

Therefore, we conclude that

$$\mathbf{E}_{\mathbf{b} \sim \mathcal{D}^\ell} [\mu(\mathbf{b})] \leq h \cdot \mathbf{E}_{X \sim \mathcal{D}} [X] + \frac{2(\ell - h)}{\rho} \left(\epsilon + \sqrt{\epsilon} \cdot \mathbf{E}_{X \sim \mathcal{D}} [\sqrt{X}] \right).$$

□

6.2 Necessity of Bayesian Incentive Compatibility

If we insist on ex post incentive compatibility, the known- h model will not help us overcome the limit on miner revenue for mechanisms that are universal in h in [SCW23].

Theorem 6.3. *Fix any $h \geq 1$, $d \geq c \geq 1$ and $\rho \in (0, 1)$. Given any (possibly randomized) MPC-assisted TFM that is ex post ϵ_u -UIC and ex post ϵ_s -SCP in an (h, ρ, c, d) -environment, it must be that for any $\mathbf{b} = (b_1, \dots, b_n)$ of length $n > h$,*

$$\mu(\mathbf{b}) \leq \frac{2n\epsilon}{\rho} + \frac{2\sqrt{\epsilon}}{\rho} \sum_{i=1}^n \sqrt{b_i}, \quad (28)$$

where $\epsilon = \epsilon_s + \epsilon_u$.

As a special case, for strict incentive compatibility where $\epsilon = 0$, we have that for any bid \mathbf{b} ,

$$\mu(\mathbf{b}) = 0.$$

Proof. In the proof we make use of the following lemma.

Lemma 6.4. *Fix any $h \geq 1$, $d \geq c \geq 1$ and $\rho \in (0, 1)$. Given any (possibly randomized) MPC-assisted TFM that is ex post ϵ_u -UIC and ex post ϵ_s -SCP in an (h, ρ, c, d) -environment, for any value v and any \mathbf{b} of length at least h , it must be that*

$$\mu(\mathbf{b}, v) - \mu(\mathbf{b}, 0) \leq \begin{cases} \frac{2}{\rho}(\epsilon_s + \epsilon_u), & \text{if } v \leq \epsilon_s + \epsilon_u \\ \frac{2}{\rho}(\sqrt{v(\epsilon_s + \epsilon_u)}), & \text{if } v > \epsilon_s + \epsilon_u. \end{cases} \quad (29)$$

For now, we assume that the lemma holds, and we show how the theorem follows. The proof to Lemma 6.4 appears afterward. For any $\mathbf{b} = (b_1, \dots, b_n)$ of length $n > h$, it must be that

$$\begin{aligned} \mu(\mathbf{b}) &= \mu(b_1, b_2, \dots, b_n) \\ &\leq \mu(b_1, \dots, b_{n-1}, 0) + \frac{2}{\rho}\epsilon + \frac{2}{\rho}\sqrt{b_n\epsilon} && \text{By Lemma 6.4} \\ &\leq \mu(b_1, \dots, b_{n-2}, 0, 0) + \frac{4}{\rho}\epsilon + \frac{2}{\rho}\sqrt{b_n\epsilon} + \frac{2}{\rho}\sqrt{b_{n-1}\epsilon} \\ &\leq \dots \leq \mu(0, \dots, 0) + \frac{2n\epsilon}{\rho} + \frac{2\sqrt{\epsilon}}{\rho} \sum_{i=1}^n \sqrt{b_i} \\ &\leq \frac{2n\epsilon}{\rho} + \frac{2\sqrt{\epsilon}}{\rho} \sum_{i=1}^n \sqrt{b_i}. \end{aligned}$$

□

Proof of Lemma 6.4. The proof is similar to the proof of Lemma 3.3 in [SCW23], while we consider the ex post setting here. To prove this lemma, we introduce the following notations. For the i -th user, we define $x_i(\mathbf{b}, v)$ to be the probability of bid v being confirmed, and $p_i(\mathbf{b}, v)$ to be the expected payment of bid v , when user i bids v and other users bid \mathbf{b} .

Lemma 6.5. *Fix any $h \geq 1$, $d \geq c \geq 1$ and $\rho \in (0, 1)$. Given any (possibly randomized) MPC-assisted TFM that is ex post ϵ_u -UIC and ex post ϵ_s -SCP in an (h, ρ, c, d) -environment, it must be that for any bid vector \mathbf{b} of length at least h , for any user i , for any $y \leq z$,*

$$\begin{aligned} & z \cdot [x_i(\mathbf{b}, z) - x_i(\mathbf{b}, y)] + \epsilon \\ & \geq p_i(\mathbf{b}, z) - p_i(\mathbf{b}, y) \\ & \geq y \cdot [x_i(\mathbf{b}, z) - x_i(\mathbf{b}, y)] - \epsilon. \end{aligned} \tag{30}$$

Proof. The proof is similar to the proof of Myerson's Lemma. For any bid vector \mathbf{b} submitted by all users other than user i , user i 's utility is $v \cdot x_i(\mathbf{b}, r) - p_i(\mathbf{b}, r)$ if its true value is v and its bid is r . Because the mechanism is ϵ_u -UIC in an (h, ρ, c, d) -environment, for all \mathbf{b} such that $|\mathbf{b}| \geq h$, it must be that

$$z \cdot x_i(\mathbf{b}, z) - p_i(\mathbf{b}, z) + \epsilon \geq z \cdot x_i(\mathbf{b}, y) - p_i(\mathbf{b}, y).$$

Otherwise, if user i 's true value is z , bidding y can bring it strictly more than ϵ utility compared to bidding truthfully, which contradicts ϵ -UIC of the mechanism in (h, ρ, c, d) -environment. By the same reasoning, for all \mathbf{b} such that $|\mathbf{b}| \geq h$, we have

$$y \cdot x_i(\mathbf{b}, y) - p_i(\mathbf{b}, y) + \epsilon \geq y \cdot x_i(\mathbf{b}, z) - p_i(\mathbf{b}, z).$$

The lemma thus follows by combining these two inequalities. \square

Lemma 6.6. *Fix any $h \geq 1$, $d \geq c \geq 1$ and $\rho \in (0, 1)$. Given any (possibly randomized) MPC-assisted TFM that is ex post ϵ_u -UIC and ex post ϵ_s -SCP in an (h, ρ, c, d) -environment, it must be that for any user i , for any bid vector \mathbf{b} of length at least h , for any $y \leq z$,*

$$\mu(\mathbf{b}, z) - \mu(\mathbf{b}, y) \leq \frac{1}{\rho}(\epsilon_u + \epsilon_s + S_{\mathbf{b}}(y, z)), \tag{31}$$

where $S_{\mathbf{b}}(y, z) = (z - y)[x_i(\mathbf{b}, z) - x_i(\mathbf{b}, y)]$.

Proof. The utility of user i is $v \cdot x_i(\mathbf{b}, r) - p_i(\mathbf{b}, r)$ if its true value is v and it bids r , while the rest of the world bids \mathbf{b} of length at least h . Henceforth, we fix an arbitrary bid vector \mathbf{b} such that $|\mathbf{b}| \geq h$. Imagine that the user i 's true value is y . If user i overbids $z > y$ instead of its true value y , then its expected utility decreases by

$$\begin{aligned} \Delta &= y \cdot x_i(\mathbf{b}, y) - p_i(\mathbf{b}, y) - [y \cdot x_i(\mathbf{b}, z) - p_i(\mathbf{b}, z)] \\ &= -y \cdot [x_i(\mathbf{b}, z) - x_i(\mathbf{b}, y)] + (p_i(\mathbf{b}, z) - p_i(\mathbf{b}, y)) \\ &\leq -y \cdot [x_i(\mathbf{b}, z) - x_i(\mathbf{b}, y)] + z \cdot [x_i(\mathbf{b}, z) - x_i(\mathbf{b}, y)] + \epsilon_u && \text{By Lemma 6.5} \\ &= (z - y) \cdot [x_i(\mathbf{b}, z) - x_i(\mathbf{b}, y)] + \epsilon_u = S_{\mathbf{b}}(y, z) + \epsilon_u. \end{aligned}$$

By ϵ_s -SCP, it must be that $\rho\mu(\mathbf{b}, z) - \rho\mu(\mathbf{b}, y) \leq \Delta + \epsilon_s$; otherwise, a user i with true value y can collude with a set of miners and overbid z instead of its true value y , while the rest honest users bid \mathbf{b} . This strategy involves only one colluding user and one bid controlled by the coalition, and it increases the coalition's utility by strictly more than ϵ_s compared to the honest strategy. This contradicts ϵ_s -SCP of the mechanism in the (h, ρ, c, d) -environment. \square

Now we proceed to prove Lemma 6.4. Let $\epsilon' = \epsilon_s + \epsilon_u$. For any user i , fix an arbitrary \mathbf{b} of length at least h from other users. Consider the following two cases.

- **Case 1: If $v \leq \epsilon'$.** In this case, by Lemma 6.6, we have that

$$\mu(\mathbf{b}, v) - \mu(\mathbf{b}, 0) \leq \frac{1}{\rho} (\epsilon_u + \epsilon_s + S_{\mathbf{b}}(0, v)) \leq \frac{1}{\rho} (\epsilon_u + \epsilon_s + v) \leq \frac{2\epsilon'}{\rho}.$$

- **Case 2: If $v > \epsilon'$.** We choose a sequence of points that partitions the interval $[0, v]$ as follows. Let $L = \lfloor \sqrt{\frac{v}{\epsilon'}} \rfloor$. Set $v_0 = 0$ and $v_{L+1} = v$. For $l = 1, \dots, L$, we set $v_l = l \cdot \sqrt{v\epsilon'}$. Each segment except the last one is of length $\sqrt{v\epsilon'}$, while the last one has a length no more than $\sqrt{v\epsilon'}$.

Now we proceed to bound $\mu(\mathbf{b}, v) - \mu(\mathbf{b}, 0)$. Note that

$$\begin{aligned} & \mu(\mathbf{b}, v) - \mu(\mathbf{b}, 0) \\ &= \sum_{l=0}^L [\mu(\mathbf{b}, v_{l+1}) - \mu(\mathbf{b}, v_l)] \\ &\leq \sum_{l=0}^L \frac{1}{\rho} [\epsilon' + S_{\mathbf{b}}(v_l, v_{l+1})] && \text{By Lemma 6.6} \\ &= \frac{L\epsilon'}{\rho} + \frac{1}{\rho} \sum_{l=0}^L (v_{l+1} - v_l) \cdot [x_i(\mathbf{b}, v_{l+1}) - x_i(\mathbf{b}, v_l)] \\ &\leq \frac{L\epsilon'}{\rho} + \frac{1}{\rho} \sqrt{v\epsilon'} \sum_{l=0}^L [x_i(\mathbf{b}, v_{l+1}) - x_i(\mathbf{b}, v_l)] && \text{By choice of } v_l \\ &\leq \frac{L\epsilon'}{\rho} + \frac{1}{\rho} \sqrt{v\epsilon'} && \text{By } x_i(\mathbf{b}, v) \leq 1 \end{aligned}$$

Since $L = \lfloor \sqrt{\frac{v}{\epsilon'}} \rfloor \leq \sqrt{\frac{v}{\epsilon'}}$, we have that

$$\mu(\mathbf{b}, v) - \mu(\mathbf{b}, 0) \leq \frac{2\sqrt{v\epsilon'}}{\rho}.$$

Lemma 6.4 thus follows.

6.3 Honest Majority of Bids

As mentioned in Section 1, we consider a “sufficient honesty” assumption but the precise statement of the assumption matters. In particular, had we assumed that the majority of bids are submitted by honest users (referred to as the “honest majority bids” assumption), then we would not be able to overcome the severe limitation on miner revenue. The theorem below states that under the honest majority bids assumption, we should still suffer from a constant miner revenue limitation.

Theorem 6.7. *Assuming that a majority number of bids are submitted by honest users. If a mechanism is Bayesian UIC, Bayesian MIC, and Bayesian SCP (even for $c = 1$) under the “honest majority bids” assumption, it must be that $\mathbf{E}_{\mathbf{b} \sim \mathcal{D}^n} [\mu(\mathbf{b})] \leq 2\mathbf{E}[\mathcal{D}]$.*

Proof. The proof is based on the following lemma.

Lemma 6.8. *Assuming that a majority of bids are submitted by honest users. Suppose that a TFM is Bayesian UIC, Bayesian MIC, and Bayesian SCP (even for $c = 1$) under the “honest majority bids” assumption. Then, as long as the number of bids $n \geq 3$, it must be that $\mathbf{E}_{\mathbf{b} \sim \mathcal{D}^n} [\mu(\mathbf{b})] \leq \mathbf{E}_{\mathbf{b}' \sim \mathcal{D}^{n-1}} [\mu(\mathbf{b}')]$.*

For now assume that Lemma 6.8 holds, and we will show how the theorem follows. The proof of Lemma 6.8 appears afterwards. By induction on n , for any $n \geq 3$,

$$\mathbf{E}_{\mathbf{b} \sim \mathcal{D}^n} [\mu(\mathbf{b})] \leq \mathbf{E}_{\mathbf{b}' \sim \mathcal{D}^2} [\mu(\mathbf{b}')] \leq 2\mathbf{E}[\mathcal{D}],$$

where the last inequality follows from the fact that the miner revenue must be upper bounded by the bids. □

Proof of Lemma 6.8. Fix a mechanism that is Bayesian UIC, MIC and SCP under the “honest majority bids” assumption. We first prove that for any $n \geq 3$, for any value v , it must be that

$$\mathbf{E}_{\mathbf{b} \sim \mathcal{D}^{n-1}} [\mu(\mathbf{b}, v)] = \mathbf{E}_{\mathbf{b} \sim \mathcal{D}^{n-1}} [\mu(\mathbf{b}, 0)]. \quad (32)$$

The proof of this is similar to that of Lemma C.4 of [SCW23]. Suppose that the number of bids is some arbitrary n . [SCW23] show that if Equation (32) does not hold, that is, if the expected miner revenue changes when some specific user lowers its bid from v to 0, and moreover, assuming that the mechanism satisfies Bayesian UIC, then, it must be that there exists some v' such that a coalition involving a single user whose true value is v' and a subset of the miners can play strategically to benefit themselves (when there are $n - 1$ other honest bids). Their proof still holds here under the “honest majority bids” assumption as long as $n \geq 3$, since when $n \geq 3$, there can be at least one strategic user colluding with the miners.

Now, for any $n \geq 3$, we have

$$\begin{aligned} \mathbf{E}_{\mathbf{b} \sim \mathcal{D}^n} [\mu(\mathbf{b})] &= \int_0^{+\infty} \mathbf{E}_{\mathbf{b}' \sim \mathcal{D}^{n-1}} [\mu(\mathbf{b}', r)] f(r) dr \\ &= \int_0^{+\infty} \mathbf{E}_{\mathbf{b}' \sim \mathcal{D}^{n-1}} [\mu(\mathbf{b}', 0)] f(r) dr && \text{By Equation (32)} \\ &= \mathbf{E}_{\mathbf{b}' \sim \mathcal{D}^{n-1}} [\mu(\mathbf{b}', 0)]. \end{aligned}$$

By Bayesian MIC, it must be that $\mathbf{E}_{\mathbf{b}' \sim \mathcal{D}^{n-1}} [\mu(\mathbf{b}', 0)] \leq \mathbf{E}_{\mathbf{b}' \sim \mathcal{D}^{n-1}} [\mu(\mathbf{b}')]$. Otherwise, when there are $n - 1$ honest bids, the miners can inject a 0 and increase the miner revenue while it does not need to pay anything for injecting the 0-bid. This violates MIC. Notice that since $n - 1 \geq 2$, the miners injecting one bid does not violate the “honest majority bids” assumption. □

References

- [ACH11] Gilad Asharov, Ran Canetti, and Carmit Hazay. Towards a game theoretic view of secure computation. In *Eurocrypt*, 2011.
- [ADGH06] Ittai Abraham, Danny Dolev, Rica Gonen, and Joseph Halpern. Distributed computing meets game theory: Robust mechanisms for rational secret sharing and multiparty computation. In *PODC*, 2006.
- [AL11] Gilad Asharov and Yehuda Lindell. Utility dependence in correct and fair rational secret sharing. *Journal of Cryptology*, 24(1), 2011.
- [AS16] Noga Alon and Joel H Spencer. *The probabilistic method*. John Wiley & Sons, 2016.
- [BCD⁺] Vitalik Buterin, Eric Conner, Rick Dudley, Matthew Slipper, and Ian Norden. Ethereum improvement proposal 1559: Fee market change for eth 1.0 chain. <https://github.com/ethereum/EIPs/blob/master/EIPS/eip-1559.md>.
- [BEOS19] Soumya Basu, David A. Easley, Maureen O’Hara, and Emin Gün Sirer. Towards a functional fee market for cryptocurrencies. *CoRR*, abs/1901.06830, 2019.
- [CCWS21] Kai-Min Chung, T-H. Hubert Chan, Ting Wen, and Elaine Shi. Game-theoretic fairness meets multi-party protocols: The case of leader election. In *CRYPTO*. Springer-Verlag, 2021.
- [CGL⁺18] Kai-Min Chung, Yue Guo, Wei-Kai Lin, Rafael Pass, and Elaine Shi. Game theoretic notions of fairness in multi-party coin toss. In *TCC*, volume 11239, pages 563–596, 2018.
- [CS23] Hao Chung and Elaine Shi. Foundations of transaction fee mechanism design. In *SODA*, 2023.
- [DR07] Yevgeniy Dodis and Tal Rabin. Cryptography and game theory. In *AGT*, 2007.
- [EFW22] Meryem Essaidi, Matheus V. X. Ferreira, and S. Matthew Weinberg. Credible, strategyproof, optimal, and bounded expected-round single-item auctions for all distributions. In Mark Braverman, editor, *13th Innovations in Theoretical Computer Science Conference, ITCS 2022, January 31 - February 3, 2022, Berkeley, CA, USA*, volume 215 of *LIPIcs*, pages 66:1–66:19, 2022.
- [enc] Encrypted mempools. talk by Justin Drake, Ethereum Foundation, <https://www.youtube.com/watch?v=XRMOCpGY3sw>.
- [FMPS21] Matheus V. X. Ferreira, Daniel J. Moroz, David C. Parkes, and Mitchell Stern. Dynamic posted-price mechanisms for the blockchain transaction-fee market. *CoRR*, abs/2103.14144, 2021.
- [FW20] Matheus V. X. Ferreira and S. Matthew Weinberg. Credible, truthful, and two-round (optimal) auctions via cryptographic commitments. In Péter Biró, Jason D. Hartline, Michael Ostrovsky, and Ariel D. Procaccia, editors, *EC ’20: The 21st ACM Conference on Economics and Computation, Virtual Event, Hungary, July 13-17, 2020*, pages 683–712. ACM, 2020.

- [GH05] Andrew V. Goldberg and Jason D. Hartline. Collusion-resistant mechanisms for single-parameter agents. In *SODA 2005*, pages 620–629, 2005.
- [GKM⁺13] Juan A. Garay, Jonathan Katz, Ueli Maurer, Björn Tackmann, and Vassilis Zikas. Rational protocol design: Cryptography against incentive-driven adversaries. In *FOCS*, 2013.
- [GKTZ15] Juan Garay, Jonathan Katz, Björn Tackmann, and Vassilis Zikas. How fair is your protocol? a utility-based approach to protocol optimality. In *PODC*, 2015.
- [GLR10] Ronen Gradwohl, Noam Livne, and Alon Rosen. Sequential rationality in cryptographic protocols. In *FOCS*, 2010.
- [GMW87] O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game. In *ACM symposium on Theory of computing (STOC)*, 1987.
- [GTZ15] Juan A. Garay, Björn Tackmann, and Vassilis Zikas. Fair distributed computation of reactive functions. In *DISC*, volume 9363, pages 497–512, 2015.
- [GY22] Yotam Gafni and Aviv Yaish. Greedy transaction fee mechanisms for (non-)myopic miners, 2022.
- [HT04] Joseph Halpern and Vanessa Teague. Rational secret sharing and multiparty computation. In *STOC*, 2004.
- [IML05] Sergei Izmailov, Silvio Micali, and Matt Lepinski. Rational secure computation and ideal mechanism design. In *FOCS*, 2005.
- [Kat08] Jonathan Katz. Bridging game theory and cryptography: Recent results and future directions. In *TCC*, 2008.
- [KMSW22] Ilan Komargodski, Shin’ichiro Matsuo, Elaine Shi, and Ke Wu. \log^* -round game-theoretically-fair leader election. In *CRYPTO*, 2022.
- [KN08] Gillat Kol and Moni Naor. Cryptography and game theory: Designing protocols for exchanging information. In *TCC*, 2008.
- [LSZ19] Ron Lavi, Or Sattath, and Aviv Zohar. Redesigning bitcoin’s fee market. In *The World Wide Web Conference, WWW 2019*, pages 2950–2956, 2019.
- [OPRV09] Shien Jin Ong, David C. Parkes, Alon Rosen, and Salil P. Vadhan. Fairness with an honest minority and a rational majority. In *TCC*, 2009.
- [PS17] Rafael Pass and Elaine Shi. Fruitchains: A fair blockchain. In *PODC*, 2017.
- [Rou20] Tim Roughgarden. Transaction fee mechanism design for the Ethereum blockchain: An economic analysis of EIP-1559. Manuscript, <https://timroughgarden.org/papers/eip1559.pdf>, 2020.
- [Rou21] Tim Roughgarden. Transaction fee mechanism design. In *EC*, 2021.
- [SCW23] Elaine Shi, Hao Chung, and Ke Wu. What can cryptography do for decentralized mechanism design. In *ITCS*, 2023.

- [Vic61] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 16(1):8–37, 1961.
- [WAS22] Ke Wu, Gilad Asharov, and Elaine Shi. A complete characterization of game-theoretically fair, multi-party coin toss. In *Eurocrypt*, 2022.
- [Yao] Andrew Chi-Chih Yao. An Incentive Analysis of Some Bitcoin Fee Designs (Invited Talk). In *ICALP 2020*.
- [YG97] Yi-Sheng Yu and Dun-He Gu. A note on a lower bound for the smallest singular value. *Linear algebra and its Applications*, 253(1-3):25–38, 1997.
- [ZCZ22] Zishuo Zhao, Xi Chen, and Yuan Zhou. Bayesian-nash-incentive-compatible mechanism for blockchain transaction fee allocation. <https://arxiv.org/abs/2209.13099>, 2022.

A MPC-Assisted, Parity-Based Mechanism

For convenience, we repeat the MPC-assisted, parity-based mechanism, which has been introduced in Section 2.2.

MPC-assisted, parity-based mechanism

// Let m be the median of the distribution \mathcal{D} such that $\Pr_{x \sim \mathcal{D}}[x \geq m] = 1/2$.

- All bids that are at least m get confirmed and pay m .
- If the number of confirmed bids is odd, then the total miner revenue is m ; else the total miner revenue is 0.

Theorem A.1. *Fix any $h \geq 1$. The MPC-assisted, parity-based mechanism satisfies ex post UIC, Bayesian MIC, and Bayesian SCP in an $(h, *, *, *)$ -environment.*

Proof. We prove the three properties separately.

UIC. Since the confirmation and the payment of each bid are independent of other bids, injecting fake bids or dropping out does not help to increase each user's utility. For a fixed user i , suppose its true value is v . When it bids b , its expected utility is $v - m$ if $b \geq m$. By direct calculation, the expected utility is maximized when $b = v$ no matter what other bids are. Thus, the MPC-assisted, parity-based mechanism satisfies UIC.

MIC. MIC follows directly from the fact that as long as there is one honest bid sampled from \mathcal{D} , the expectation of the parity is $\frac{1}{2}$, no matter what other bids are. Therefore, the expected total miner revenue is always $\frac{1}{2}m$.

SCP. By the same reasoning as in MIC, no matter how the coalition deviates, they cannot increase the total miner revenue. Henceforth, we only need to argue that no matter how the coalition deviates, it cannot increase the joint users' utility. This follows directly from UIC and the fact that each bid is confirmed independently from other bids. \square